Netherite: Efficient Execution of Serverless Workflows

Wangyi Chen, Zhenhuan Wu

BOSTON UNIVERSITY

User Story With A Server

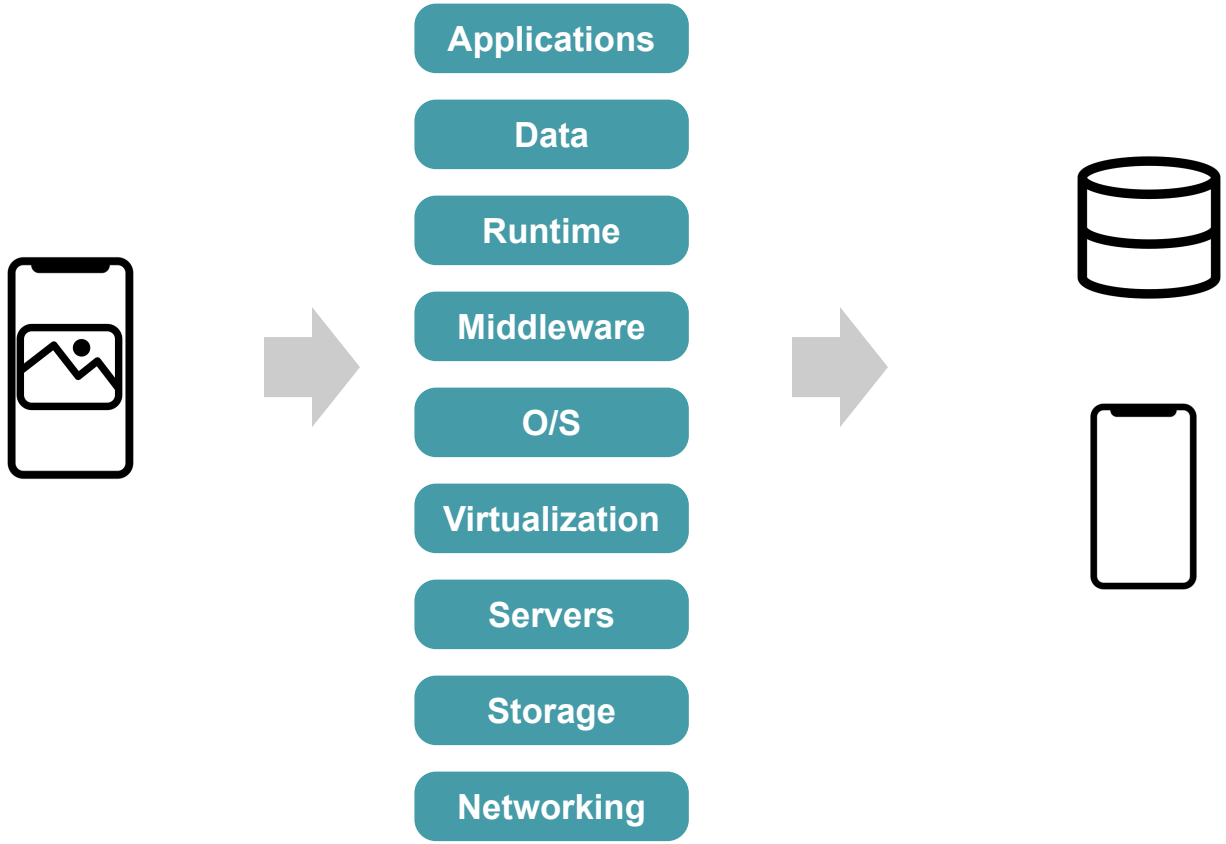

User Story With Serverless

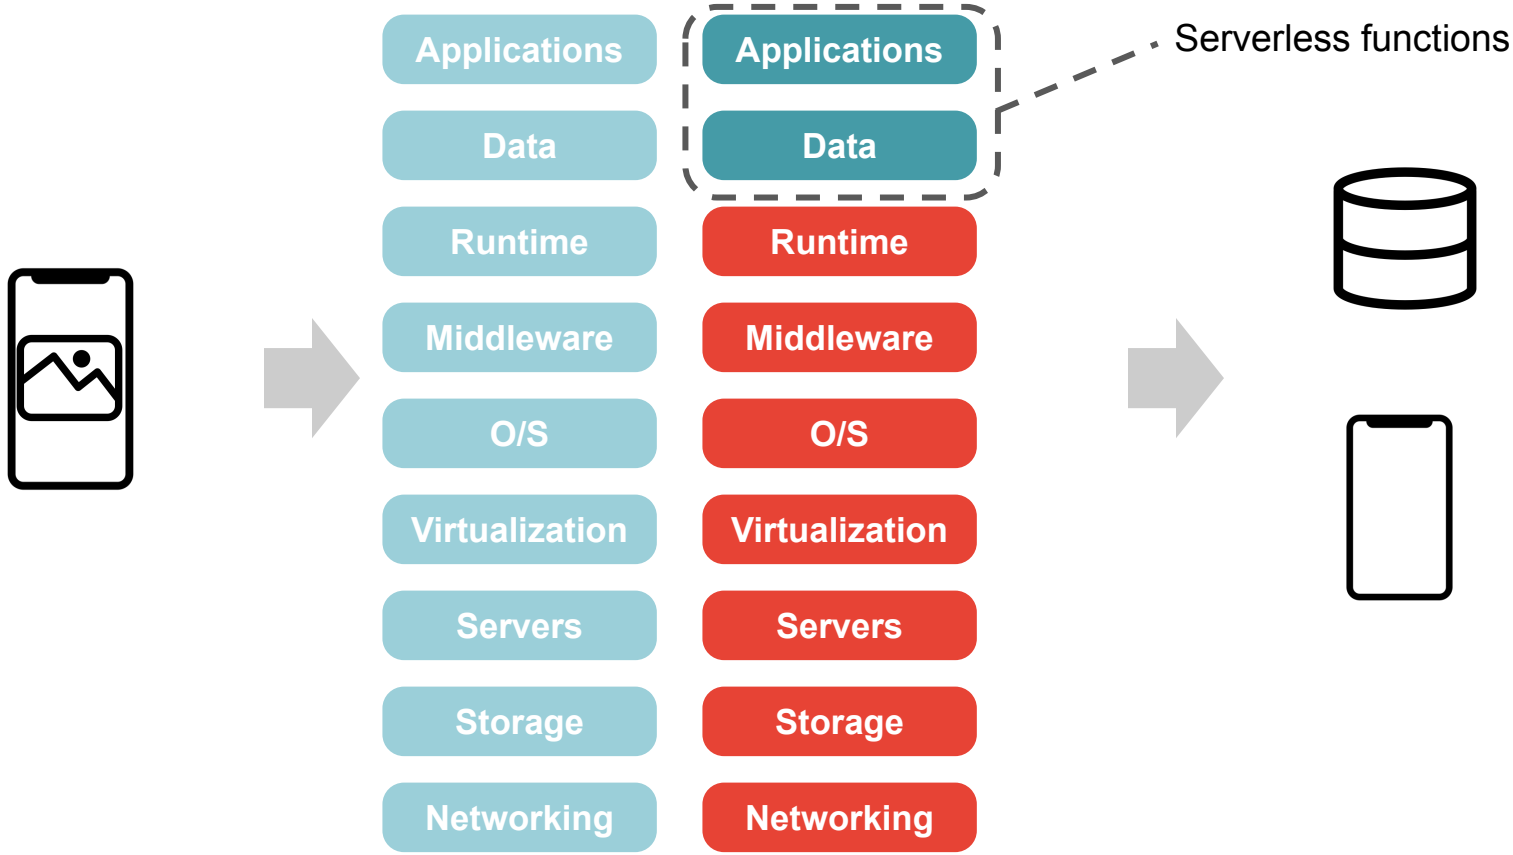

Server vs Serverless

Items	Server	Serverless
Configuration	Manual	**Automatic**
Scalability	Fixed	**Unlimited**
Elasticity	Fixed	**Elastic**
Billing	Fixed	**Pay as you use**
Load Balancing	Difficult	**Easy**

How do serverless functions work?

Serverless functions are a cloud computing service that allows developers to run code without managing infrastructure.

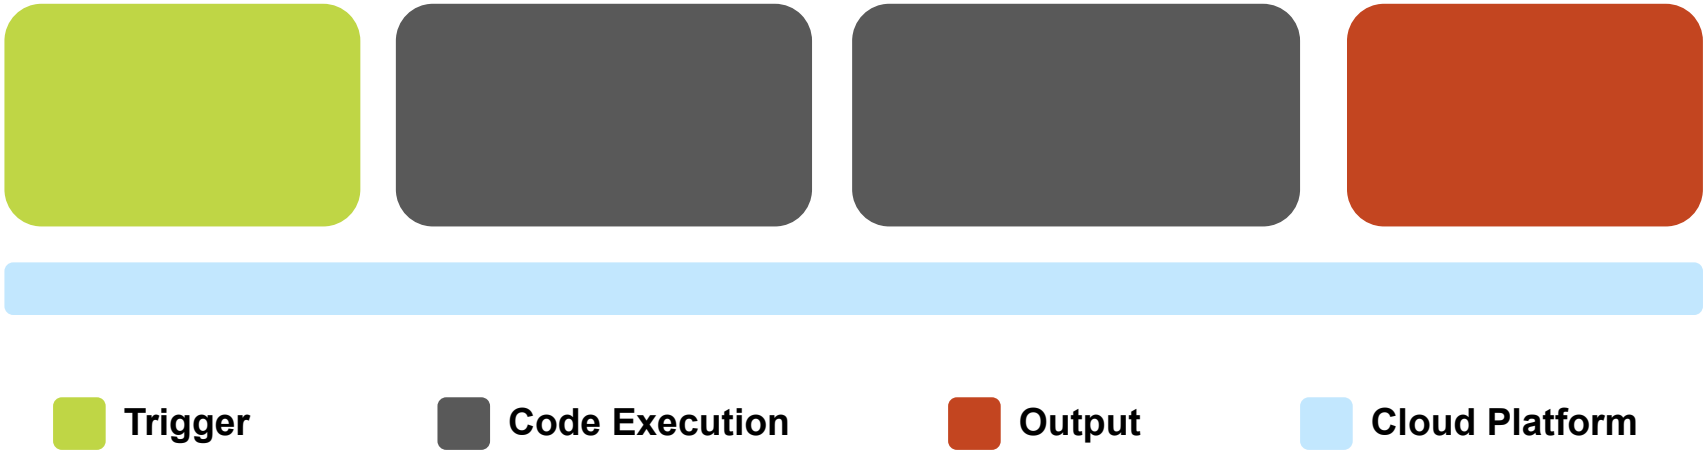

What are serverless functions?

- HTTP
- Add/update blob storage
- Add/Update database
- Scheduled task
- Process message queues

- Implement web endpoint
- Run the uploaded code
- Run custom logic

- Return to user
- Update storage
- Update database
- Return to IoT device

Executions are **STATELESS**

Trigger **Code Execution** **Output** **Cloud Platform**

What is a **STATE ?**

STATE refers to the condition of any given time, or prior knowledge of a task.

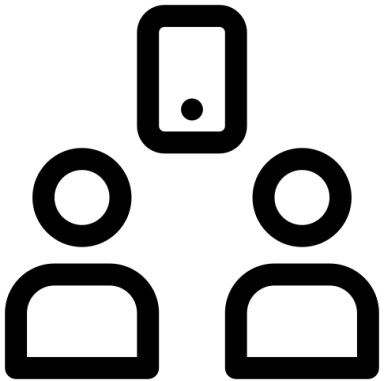

Stateful

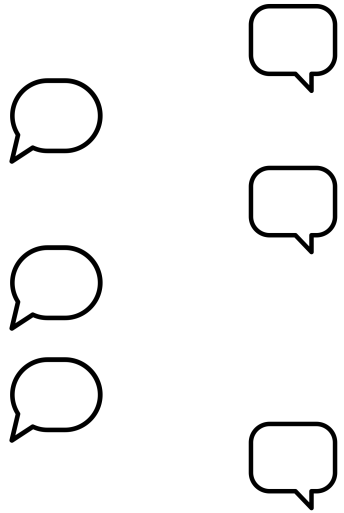

Stateless

How do stateless functions work ?

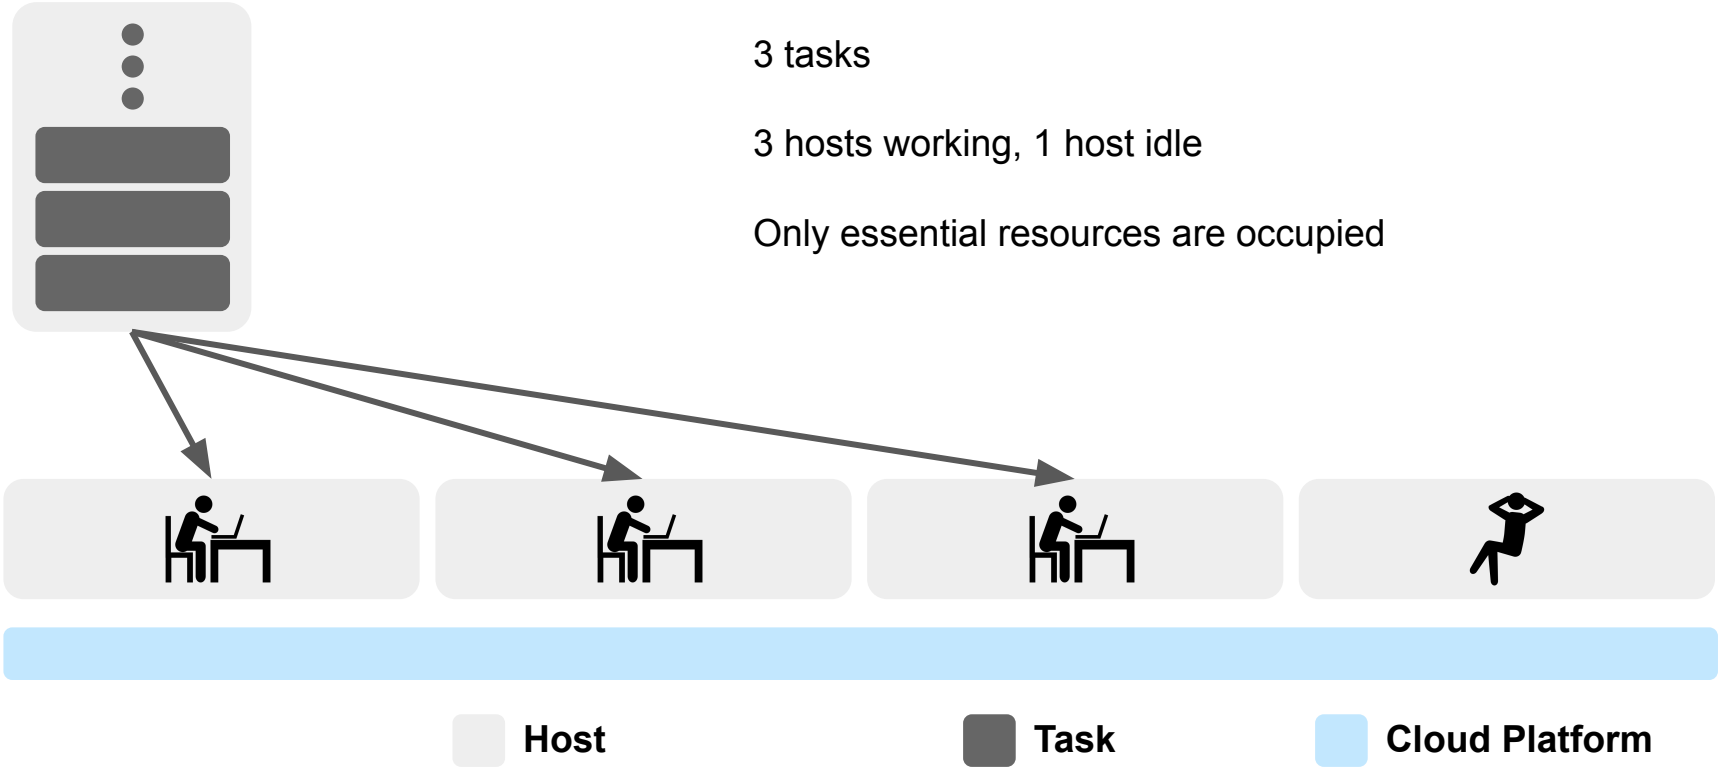

3 tasks

3 hosts working, 1 host idle

Only essential resources are occupied

History background

Stateless HTTP triggers make up less than 50% of invocations

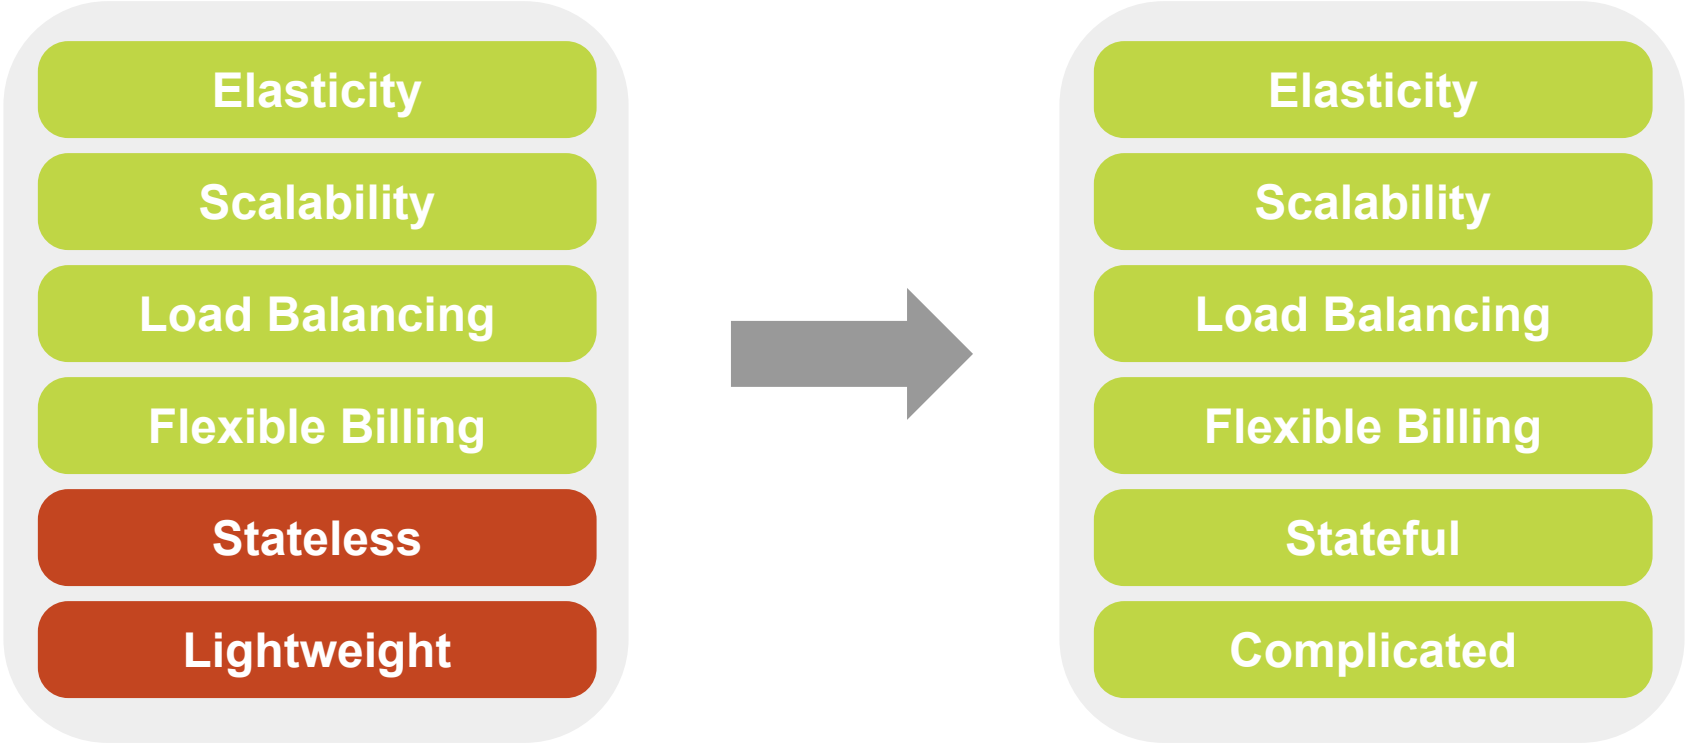

What could possibly go wrong?

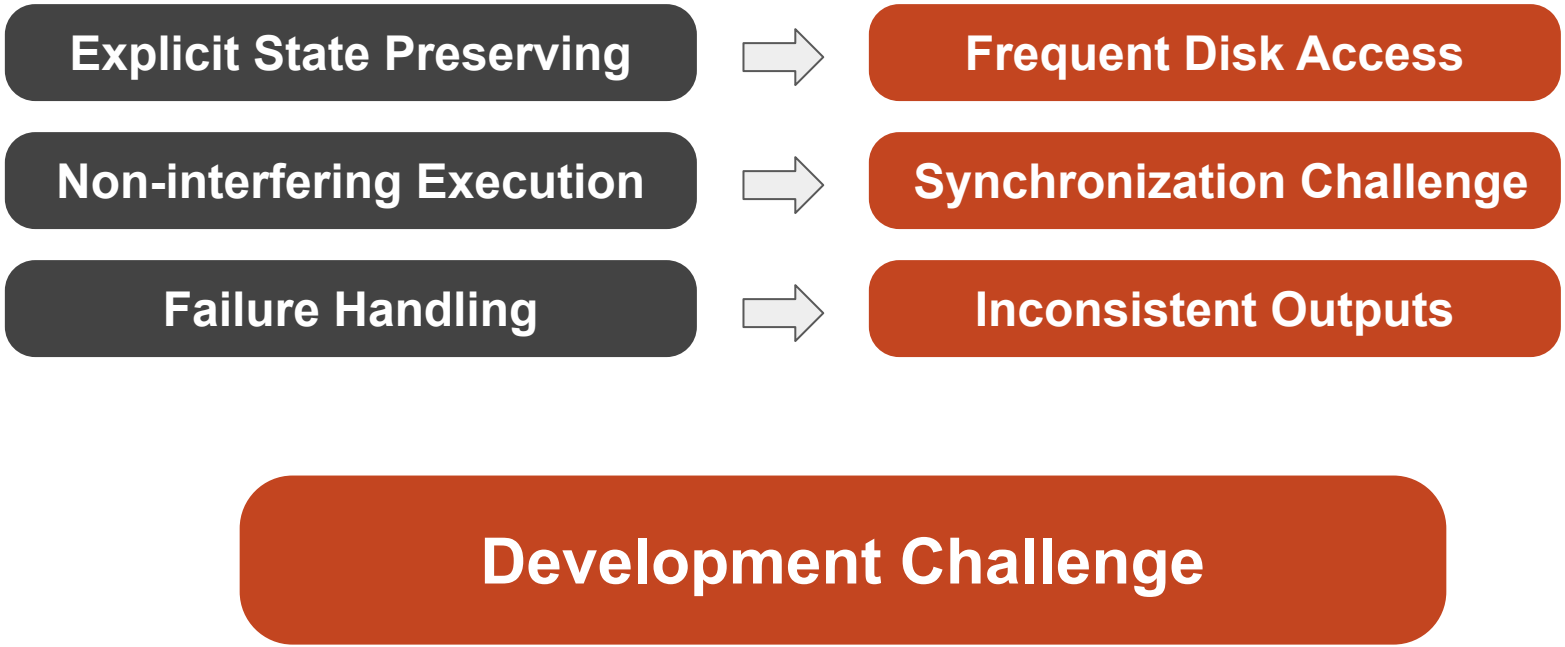

Key motivation

Less Frequent Disk Access

Implicit Synchronizations

Implicit Failure Handling

Stateful Function Structure

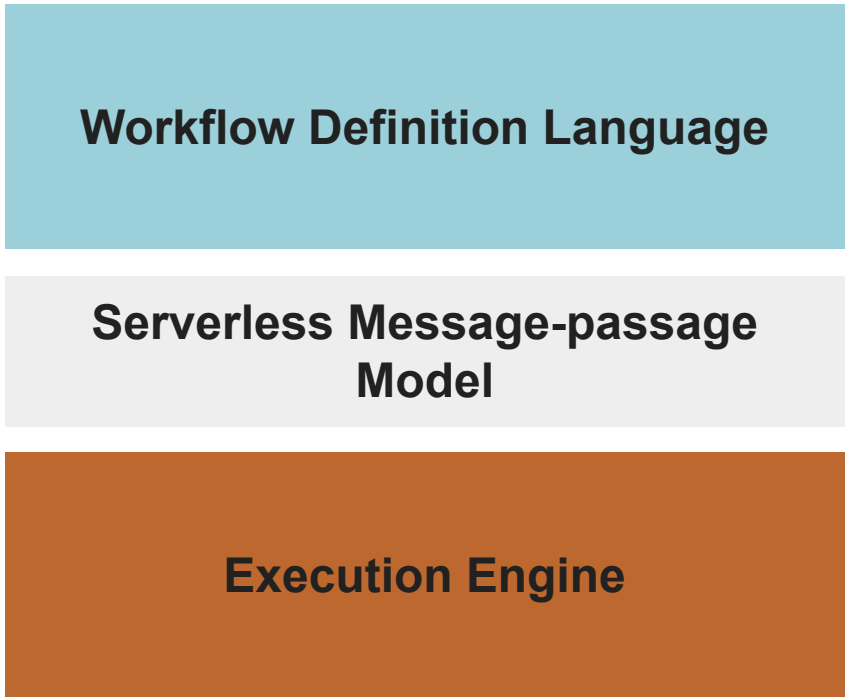

Durable functions - Azure
Step functions - AWS
…

Intermediate layer for decoupling

Azure storage
Netherite
MS SQL server

Durable Functions

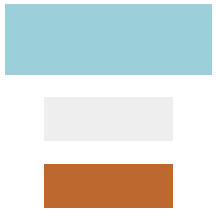

Orchestration function

Entity function

Critical Section

Orchestration Functions

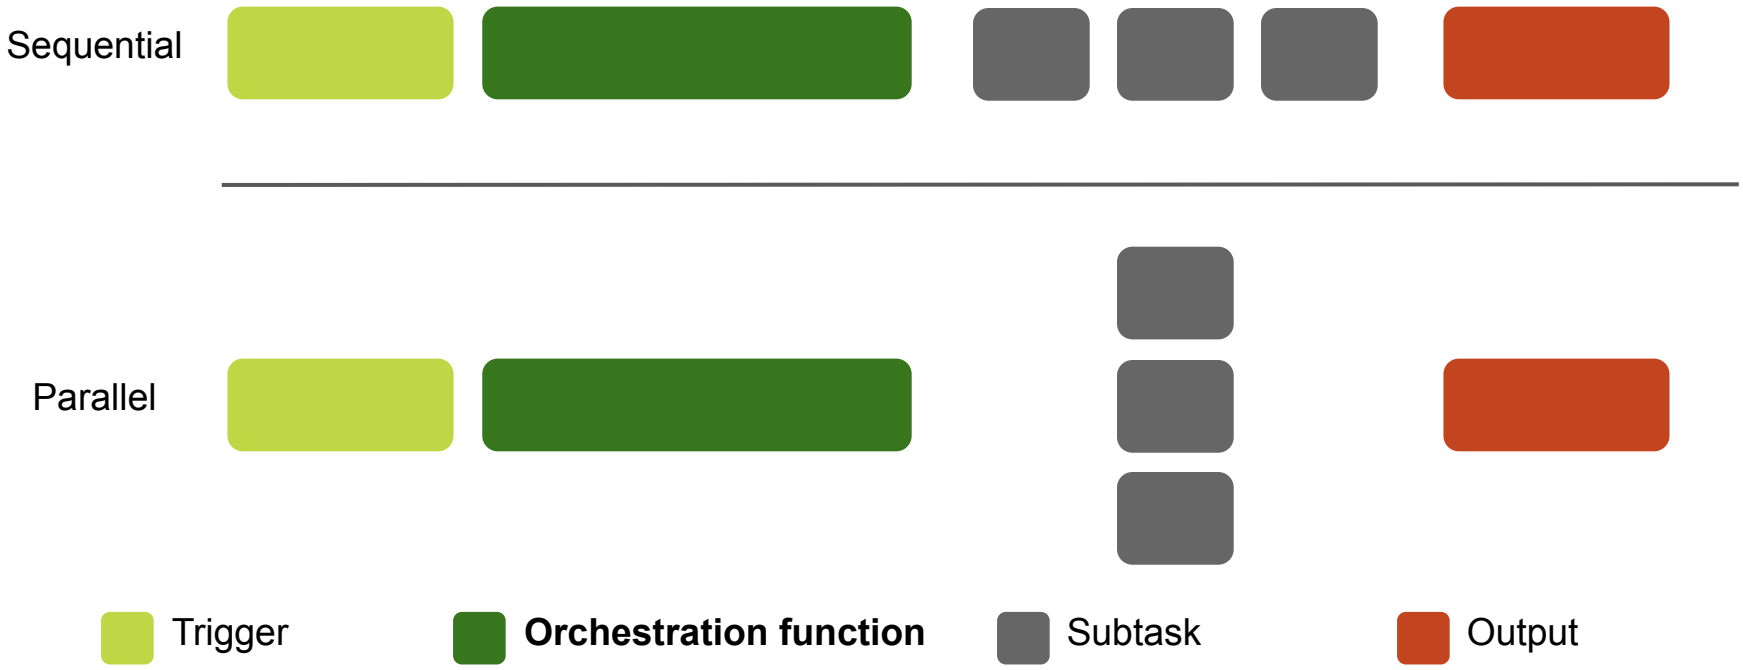

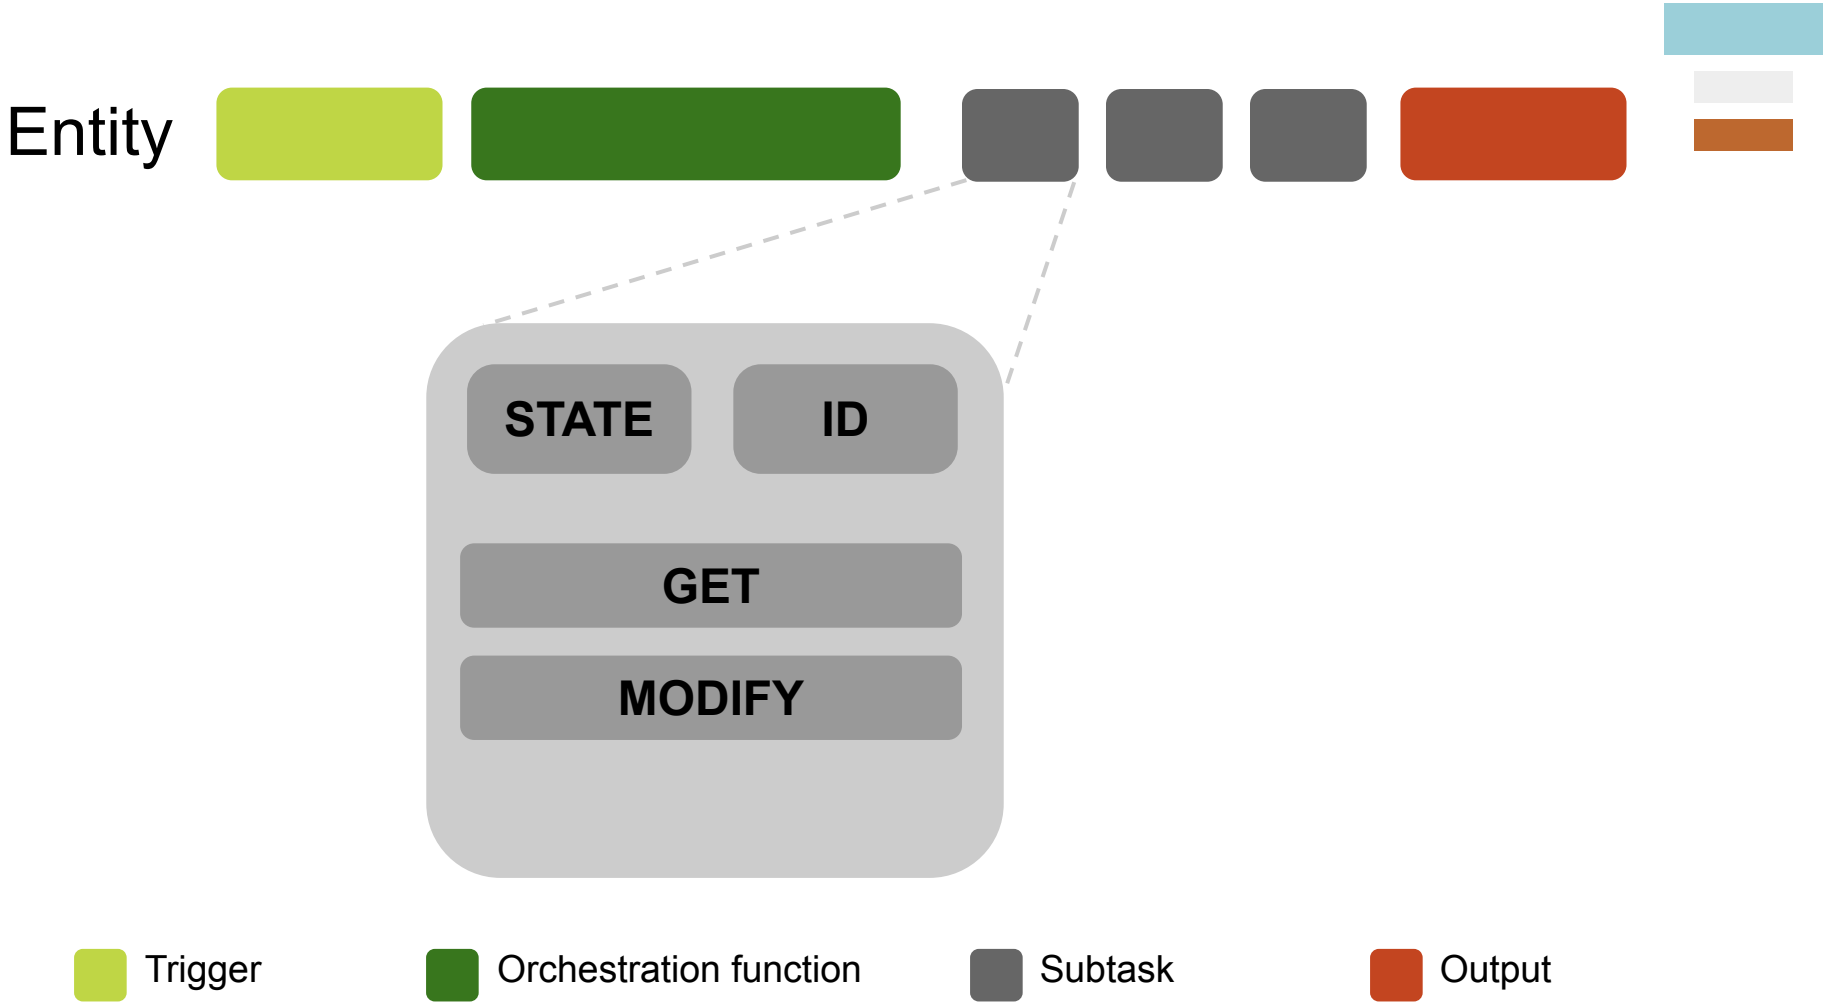
Entity
STATE
ID
GET
MODIFY
Trigger
Orchestration function
Subtask
Output

Critical Sections

Asynchronous locks

Trigger Orchestration function Entity Critical Section Output

Core Problems

Frequent Disk Access

Synchronization Challenge

Inconsistent Outputs

Message-passing Model

Stateful Function Structure

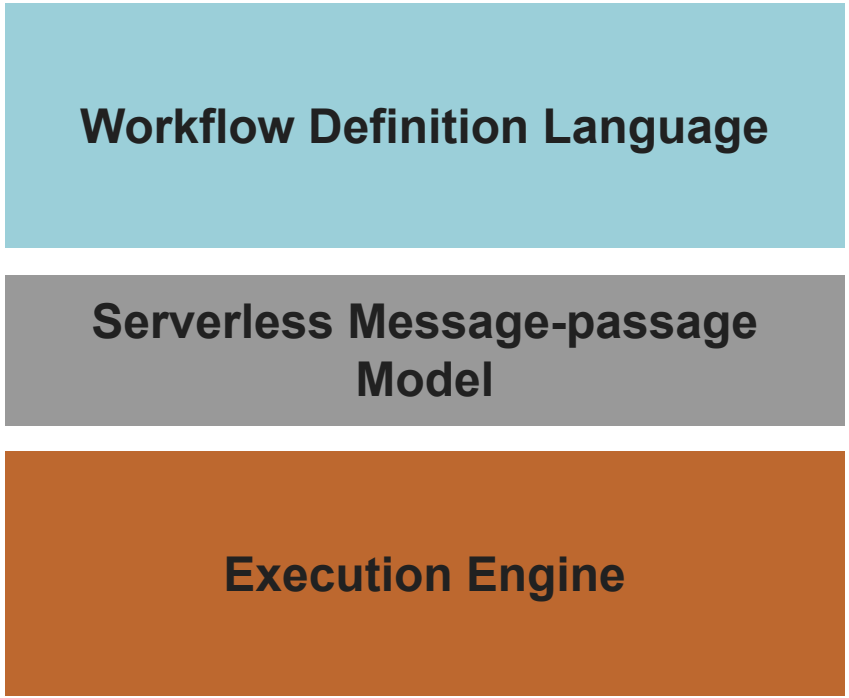

Durable functions - Azure
Step functions - AWS
…

Intermediate layer for decoupling

Azure storage
Netherite
MS SQL server

Message-passing Model Architecture

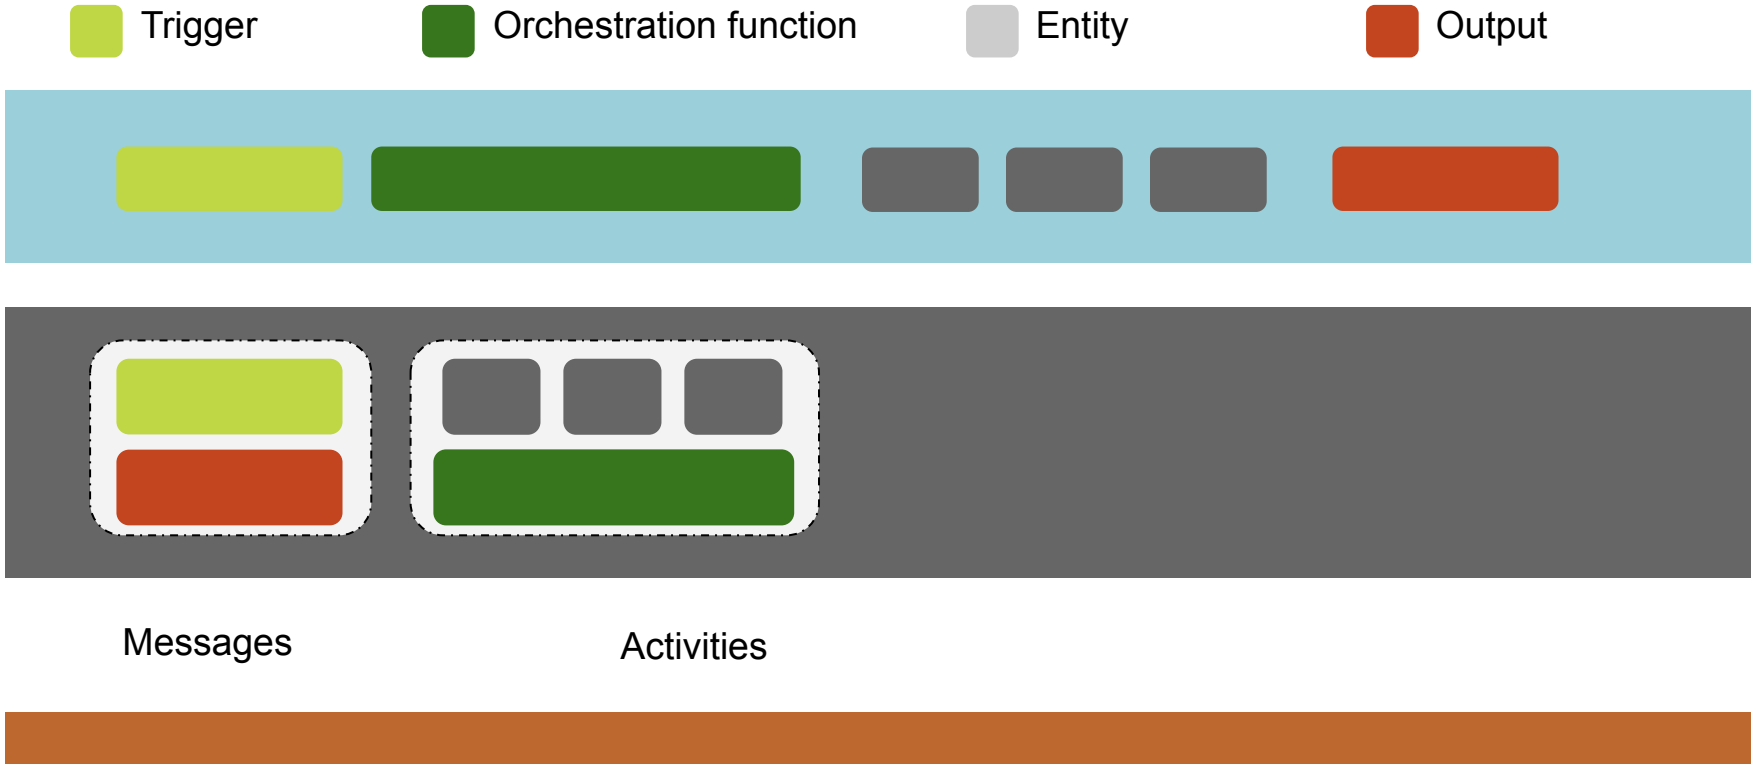

Message-passing Model Architecture

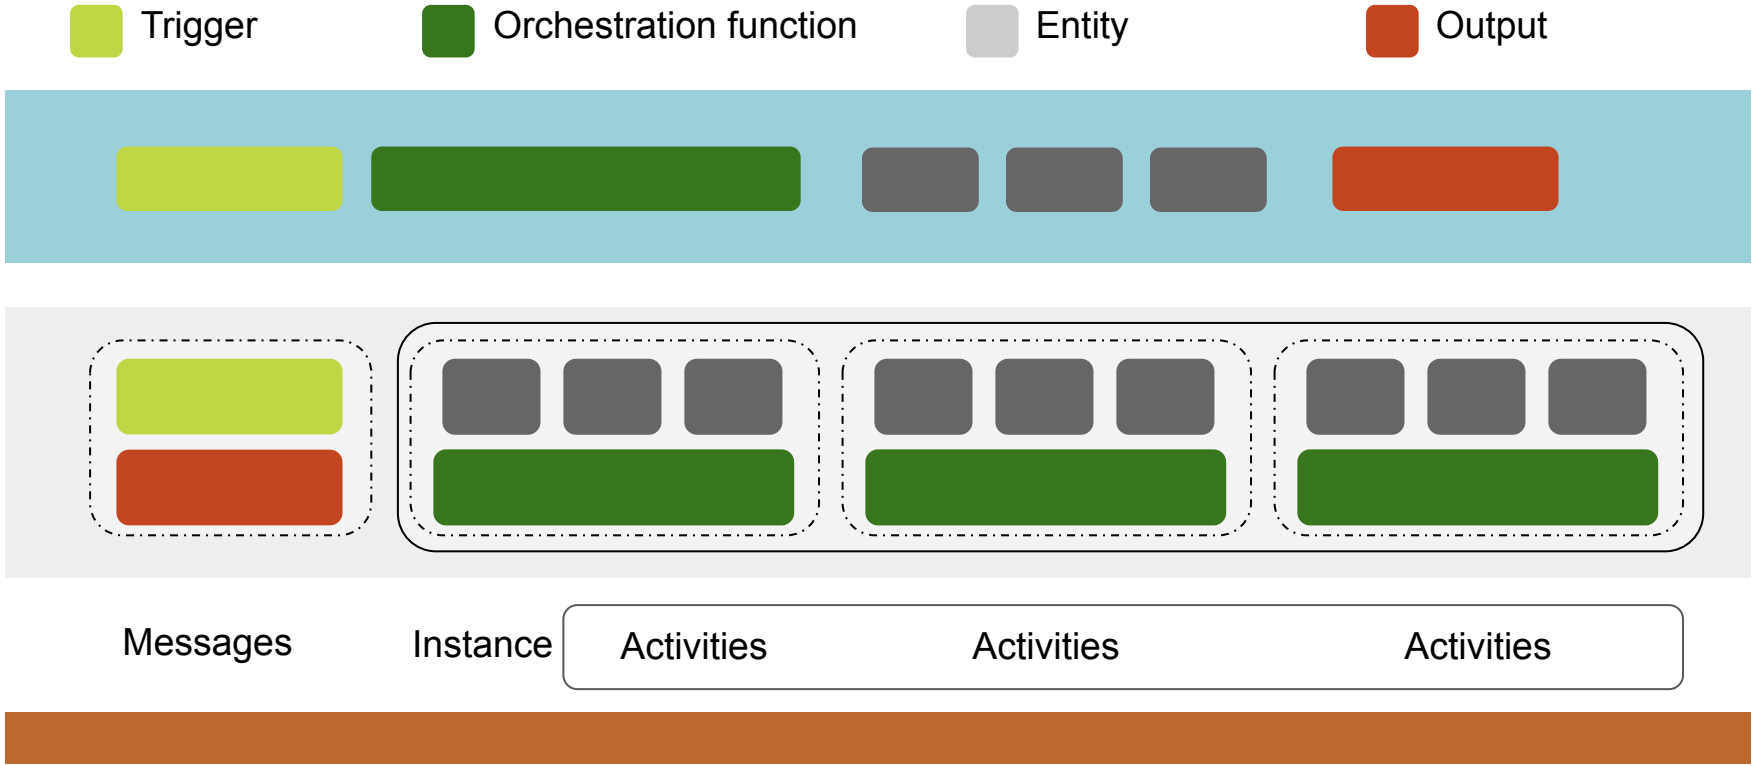

Message-passing Model Implementations

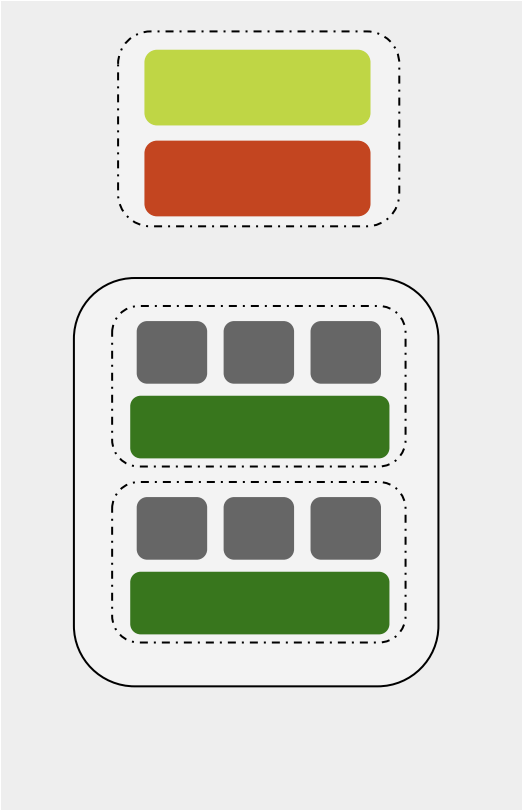

Message-passing Model

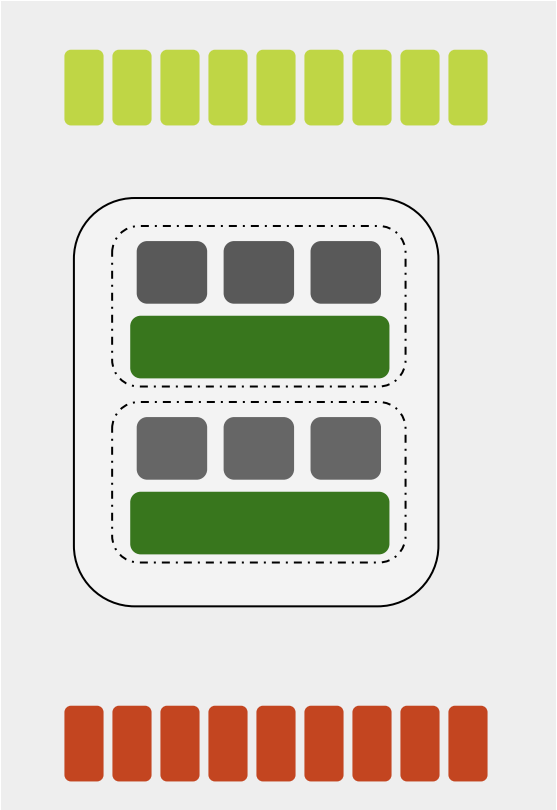

Messages (Triggers and outputs) are stacked in queues.

Message Activity Instance

Message-passing Model

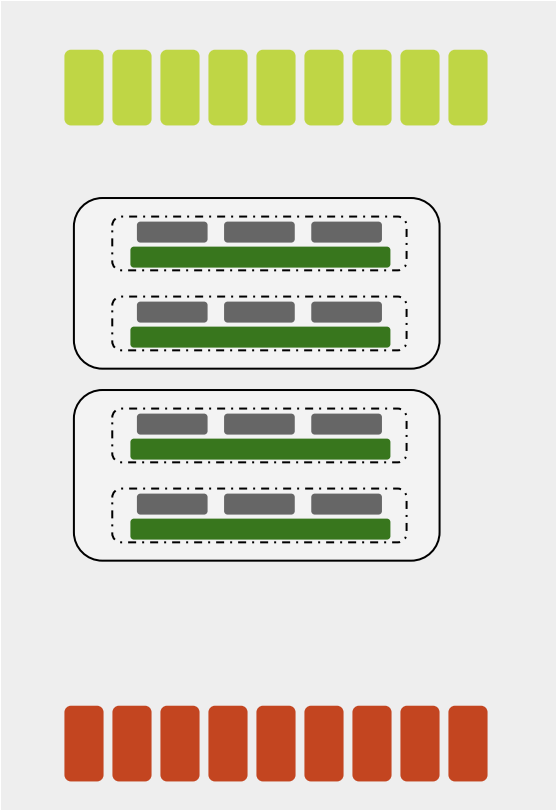

Multiple instances are running in parallel

Message Activity Instance

Message-passing Model

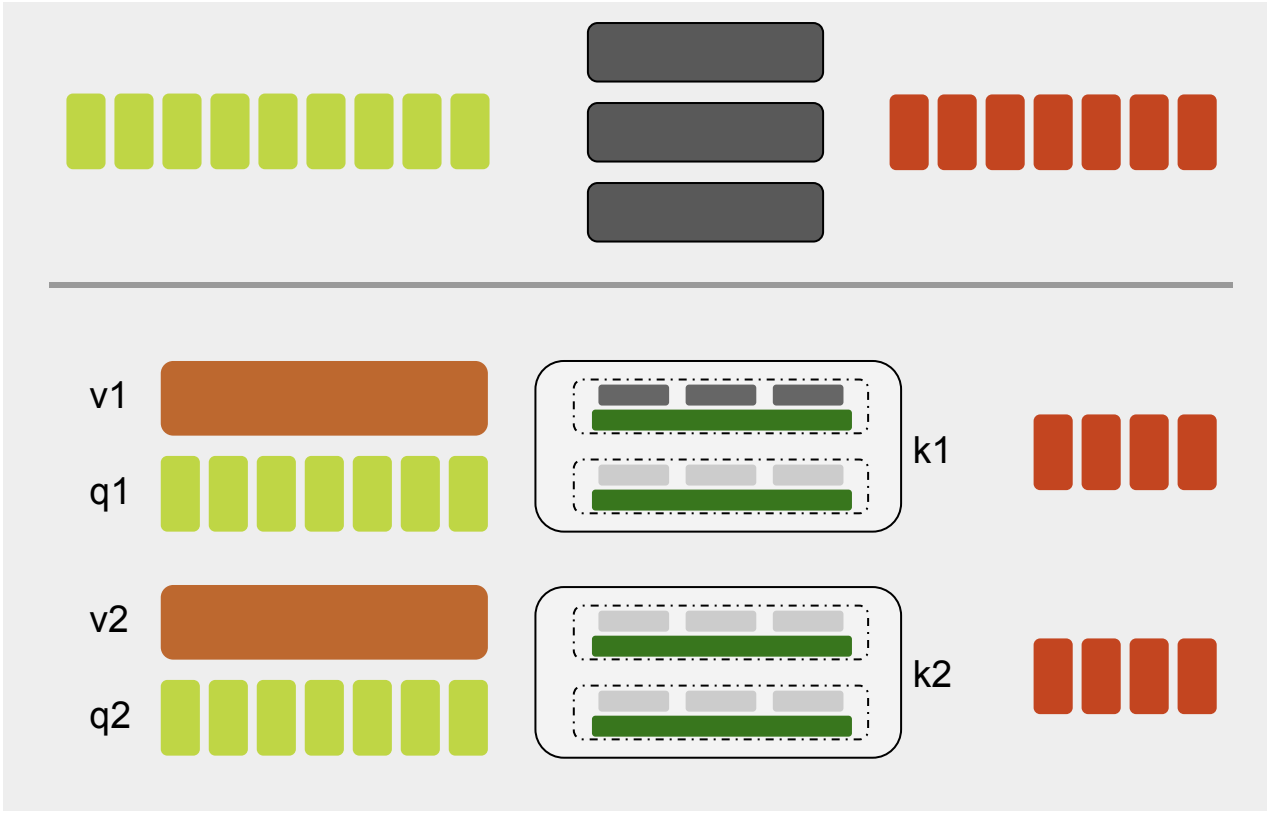

Message-passing Model

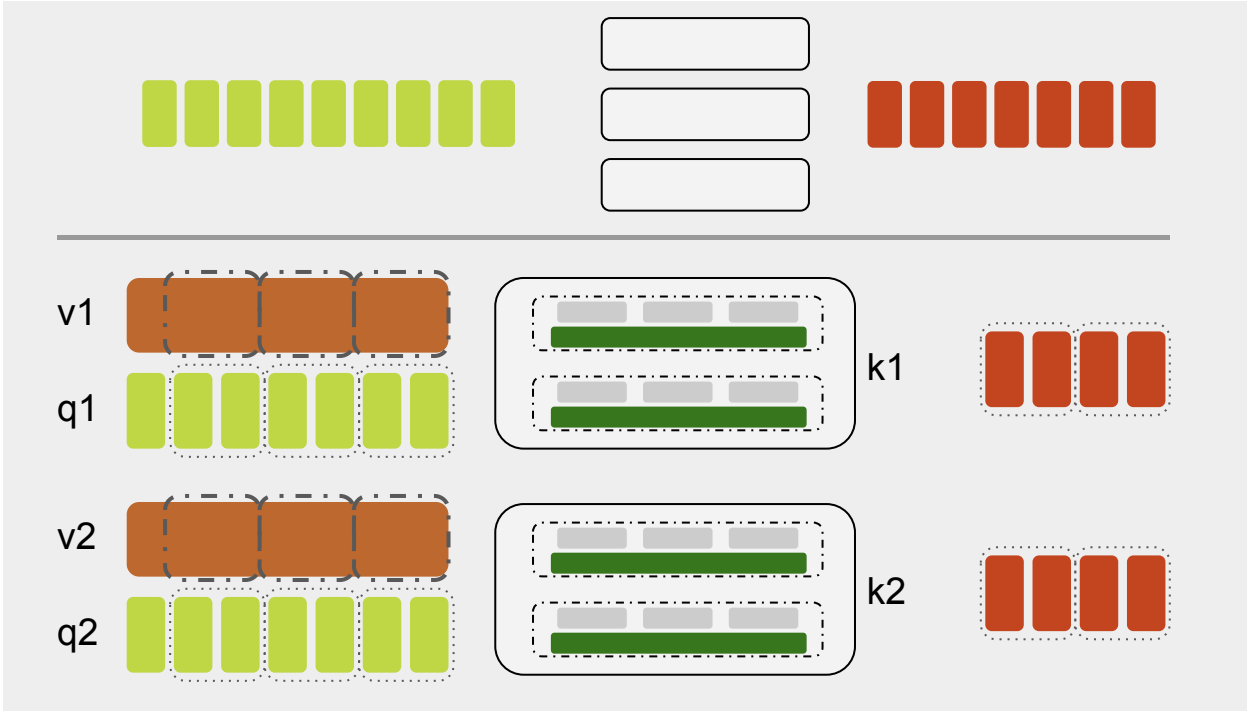

Message State Activity Instance

Frequent Disk Access

Synchronization

Inconsistent Outputs

Message-passing Model Reliability

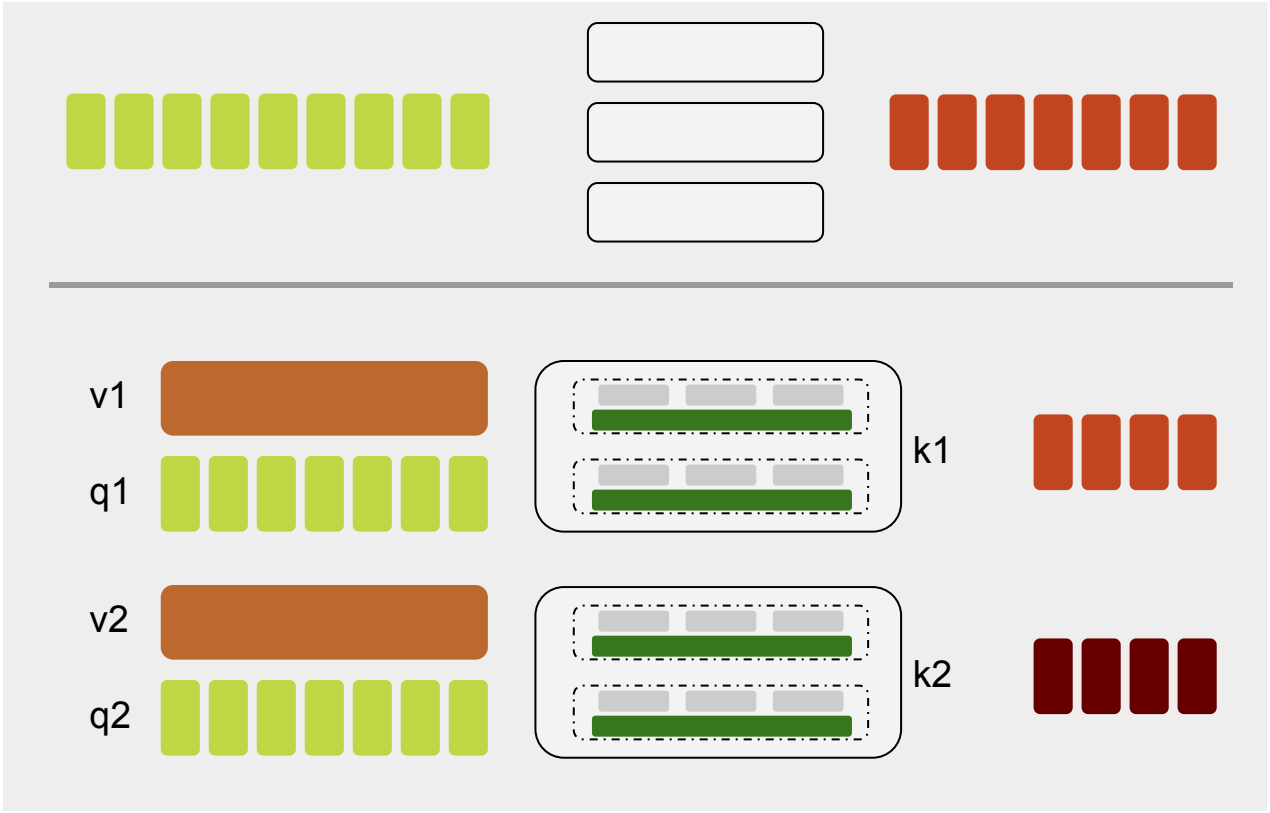

Frequent Disk Access

Synchronization

Inconsistent Outputs

Serialized consumption
Serialized commit

Message State Activity Instance

Message-passing Model Reliability

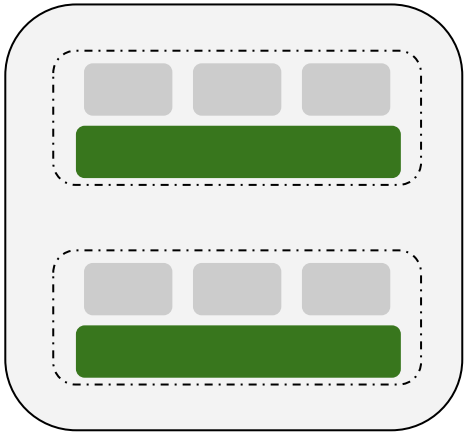

Activity Instance

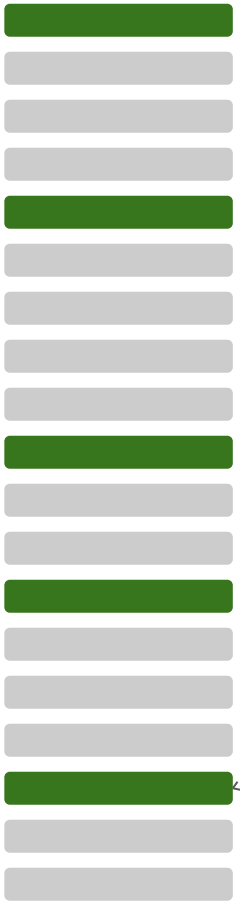

Frequent Disk Access

Synchronization

Inconsistent Outputs

Events

```
State: Fetched

Task 1: Done
Task 1 result: 22
Task 2: In progress
Task 3: pending

Output: None
```

Message-passing Model Reliability

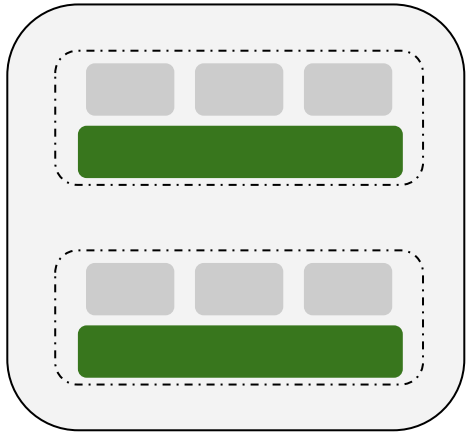

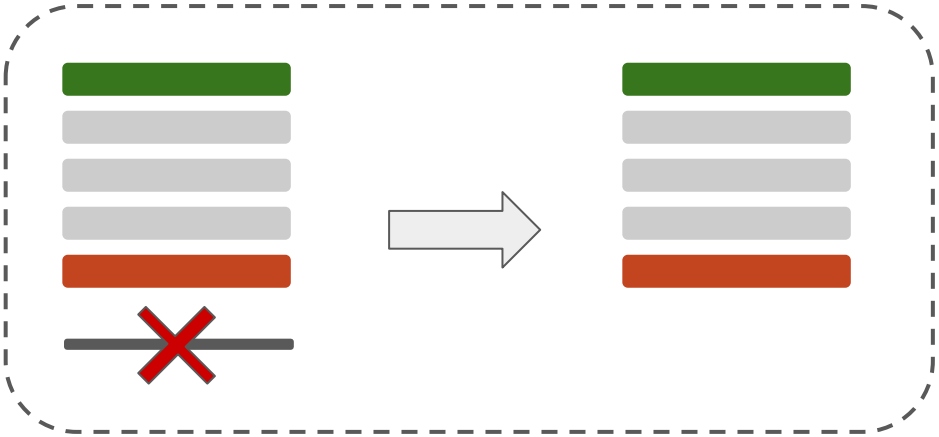

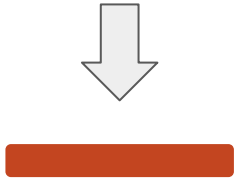

Exactly-once execution

Frequent Disk Access

Synchronization

Inconsistent Outputs

Message-passing Model Reliability

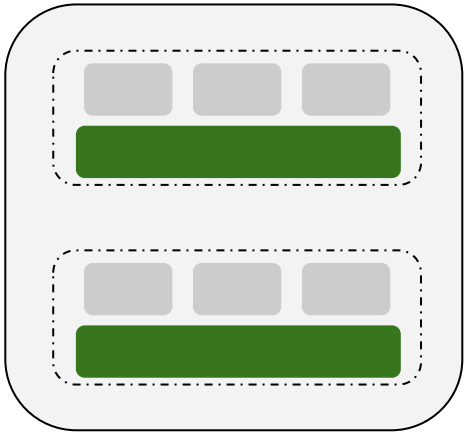

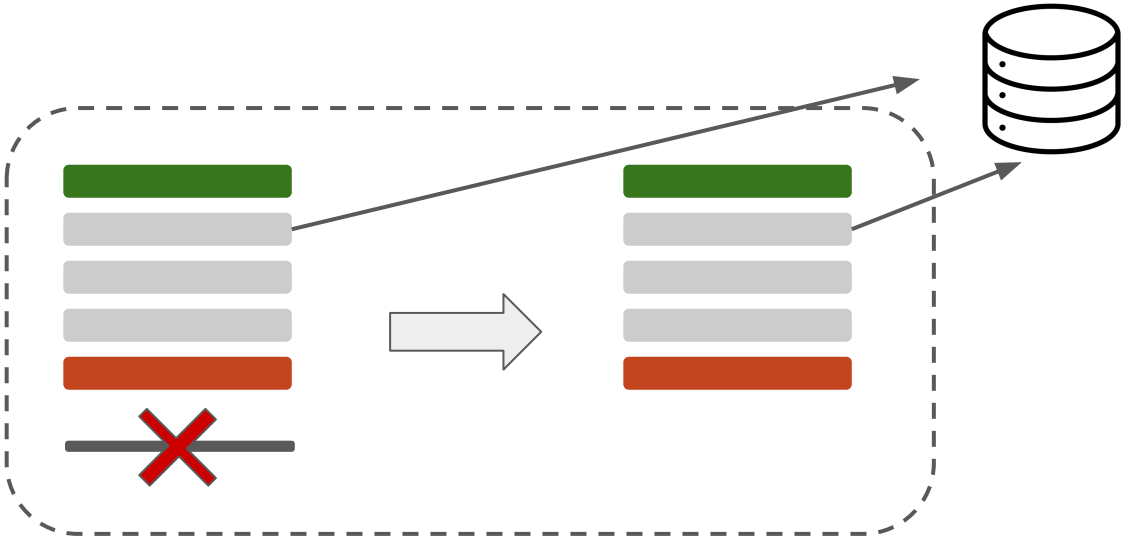

Activity Instance

There are external effects!

Frequent Disk Access

Synchronization

Inconsistent Outputs

Message-passing Model Reliability

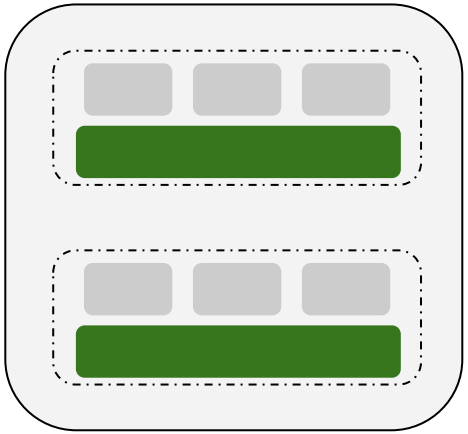

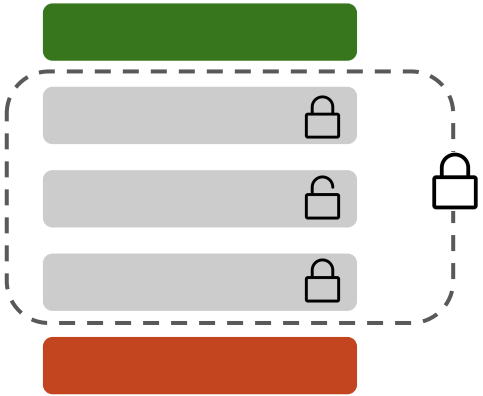

Two-phase locking protocol

Activity Instance

Frequent Disk Access

Synchronization

Inconsistent Outputs

Frequent Disk Access

Explicit Resource Isolation

Duplicate/Error Outputs

Batch Operation

Two-phase Lock Protocol

Serialized Commit

Partial History

Rollback

External Effects

Execution Engines

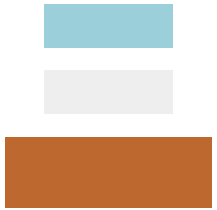

Original Design

Netherite

Stateful Function Structure

Workflow Definition Language

Durable functions
Step functions
…

Serverless Message-passage Model

Intermediate layer for decoupling

Execution Engine

Azure storage
Netherite
MS SQL server

Original Design Workflow

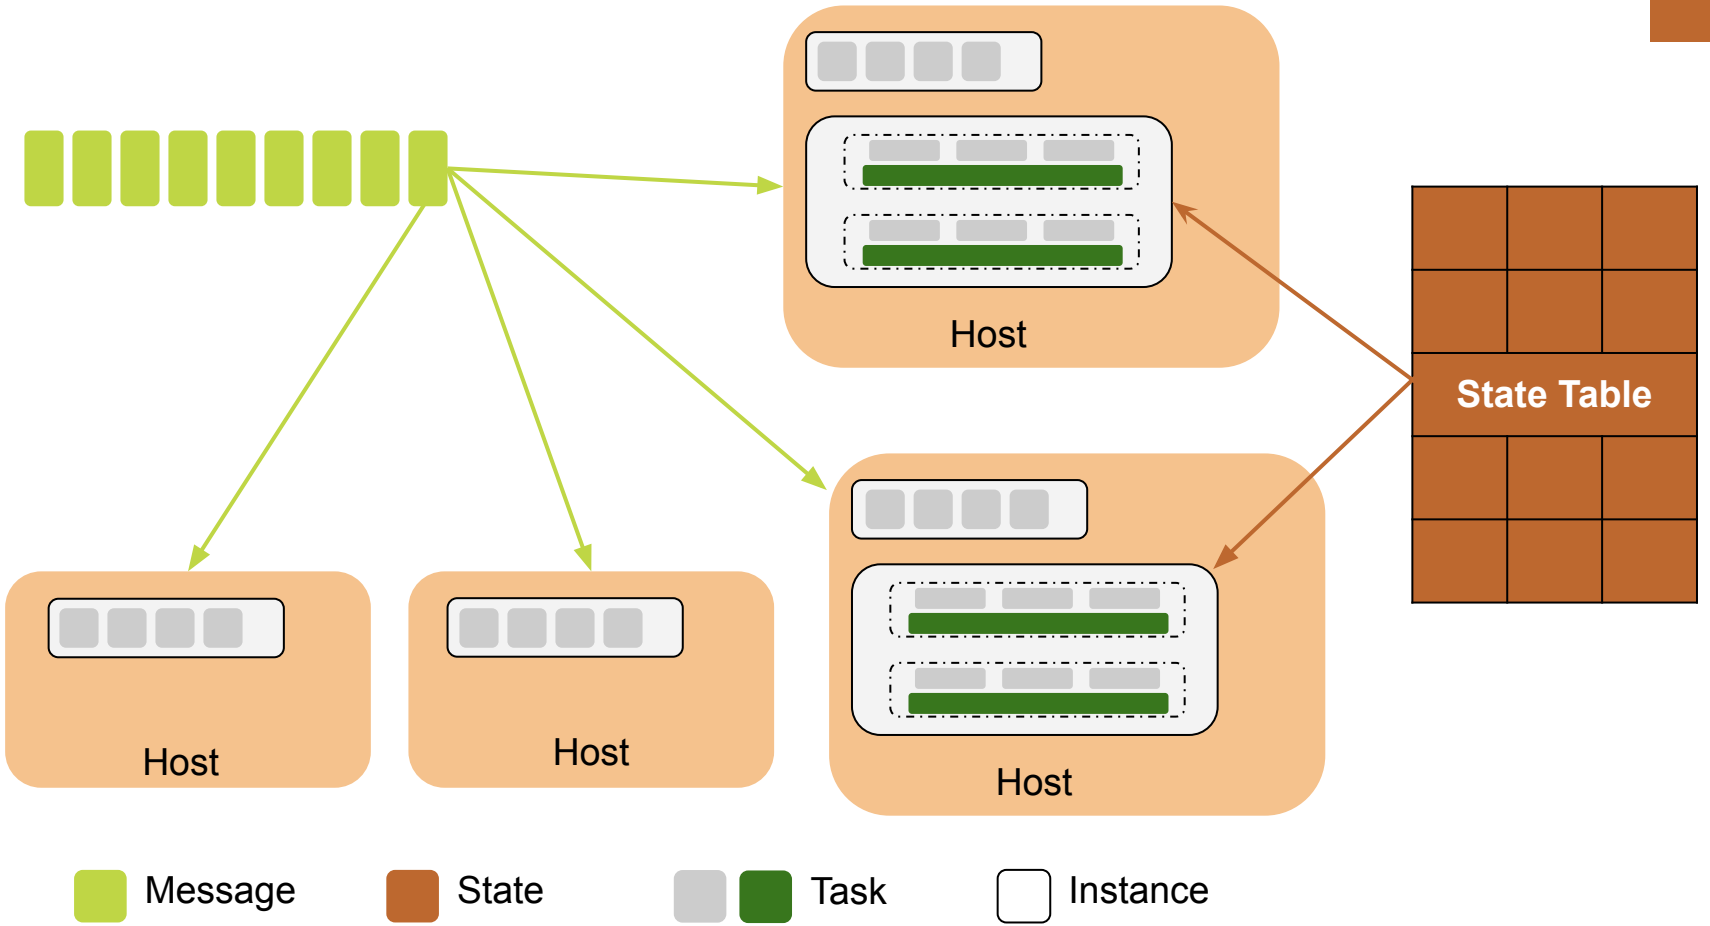

Performance and Cost

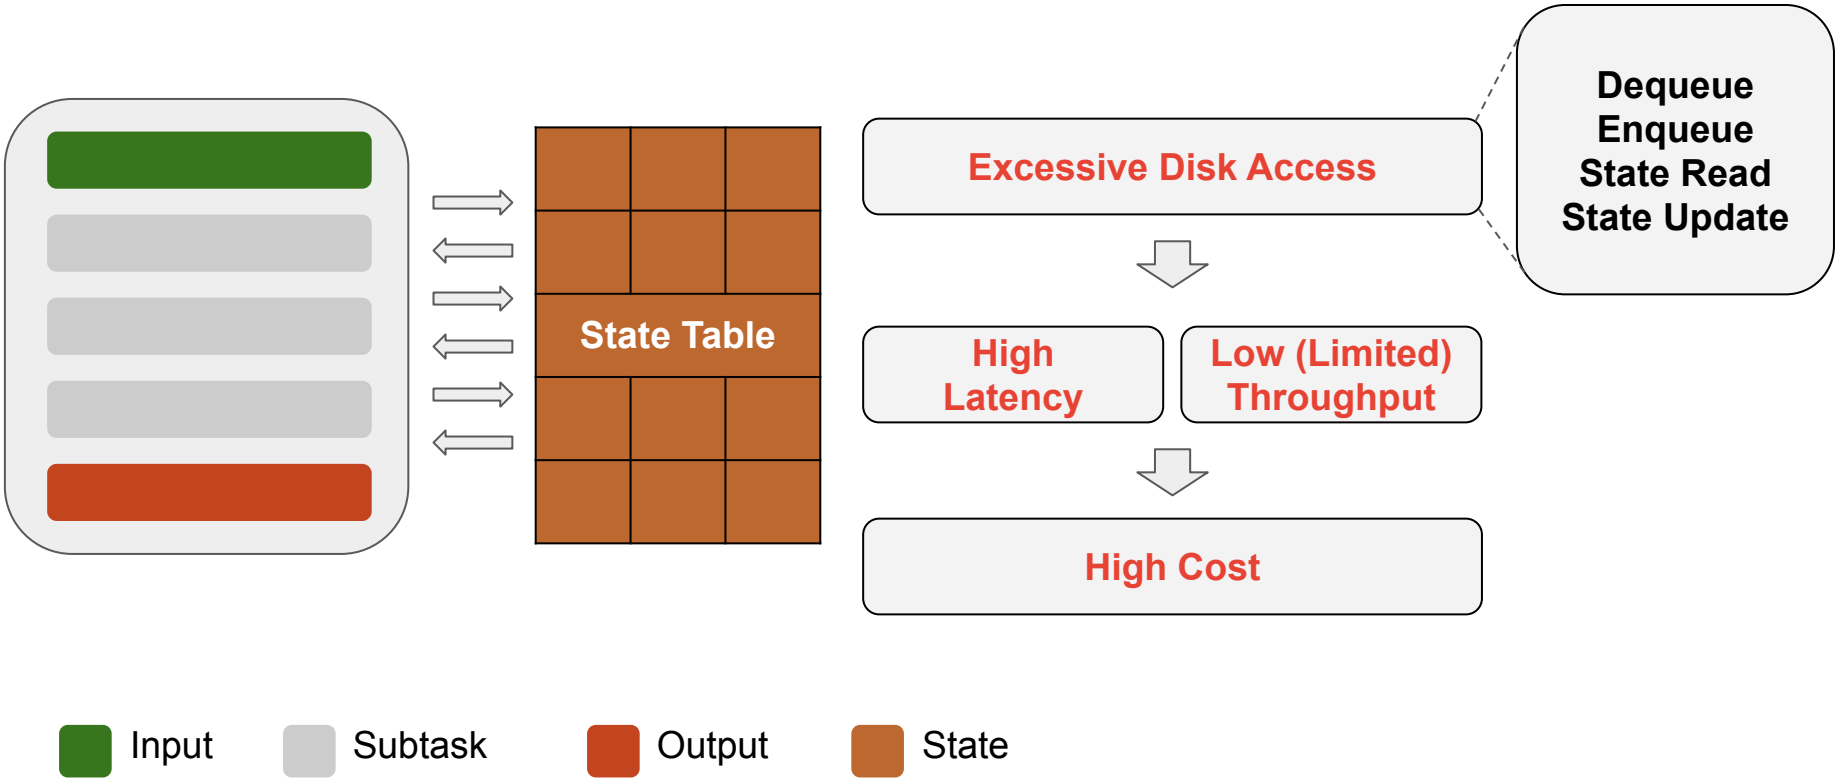

Core Problems

Frequent Disk Access

Synchronization Challenge

Inconsistent Outputs

Execution Engines

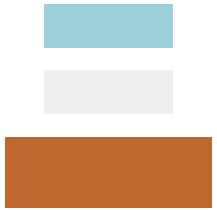

Original Design

Netherite

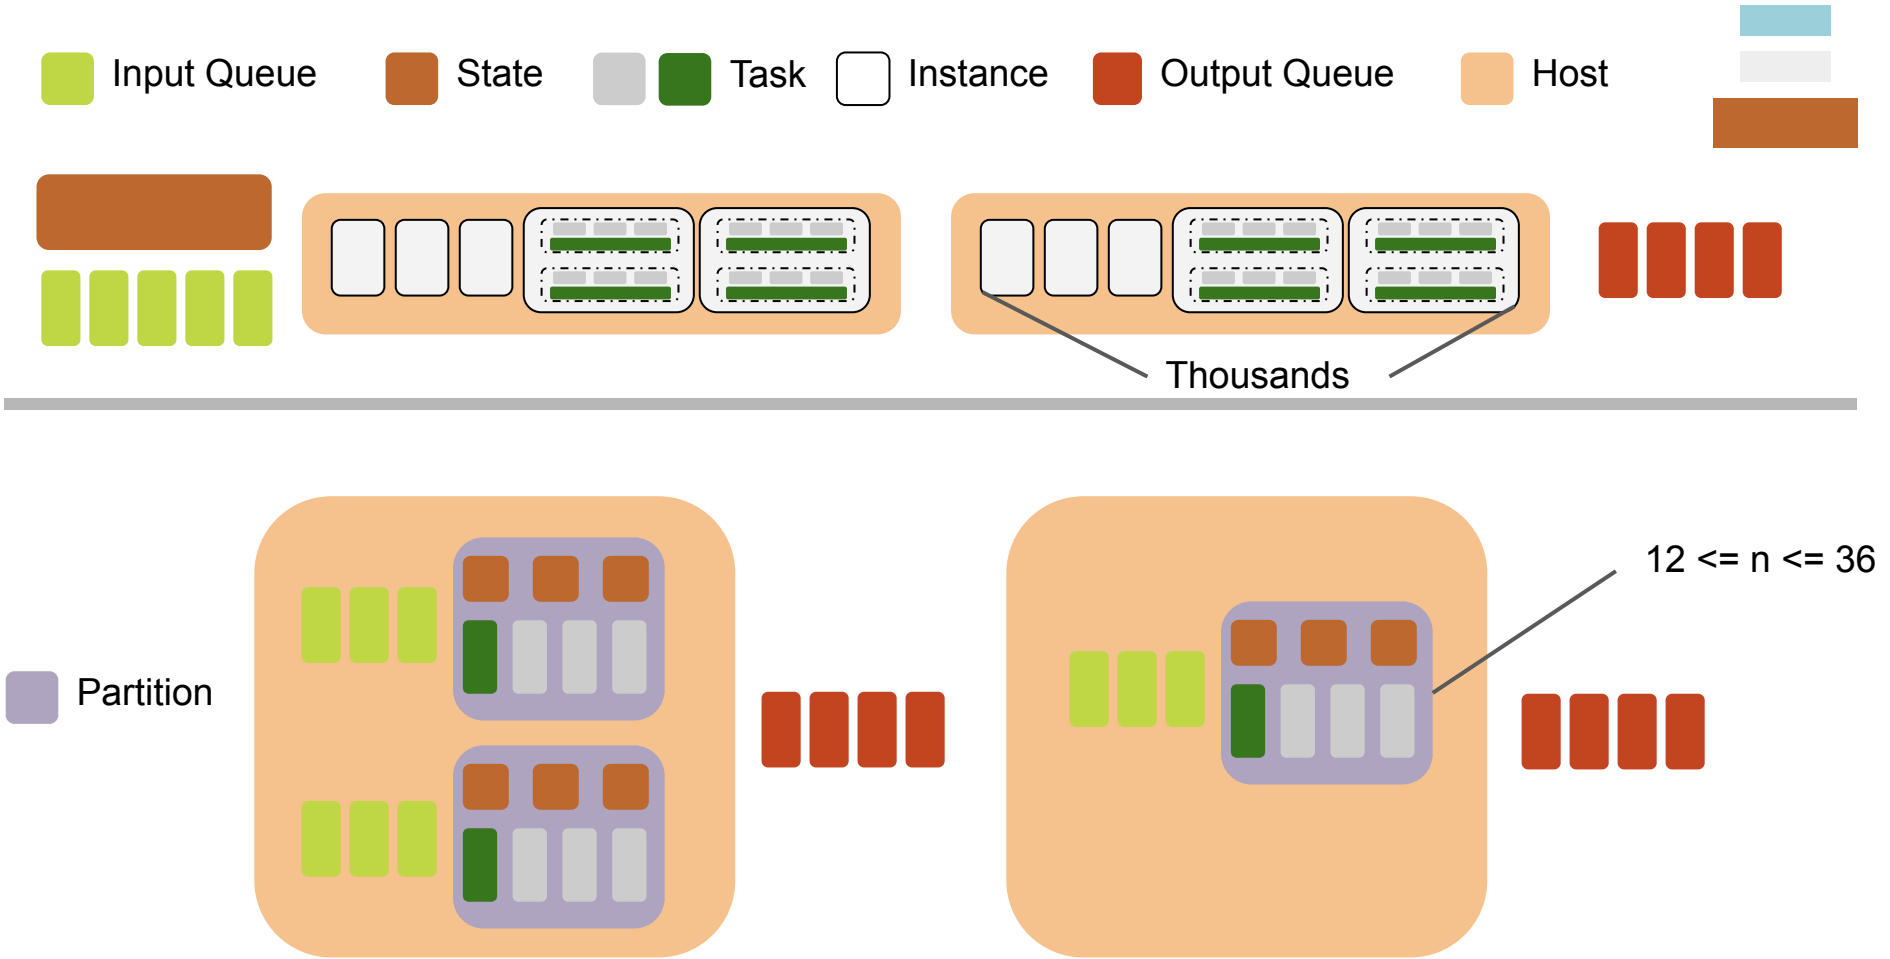
Input Queue
State
Task
Instance
Output Queue
Host
Thousands
Partition
12 <= n <= 36

Netherite - Load Balancing

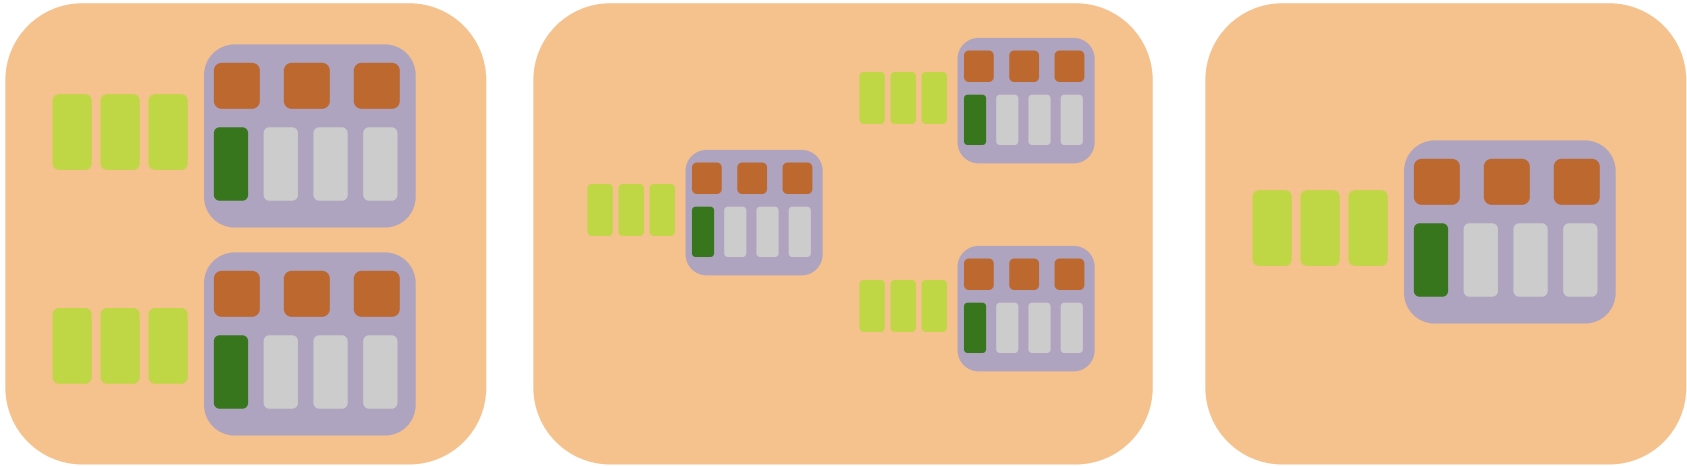

Frequent Disk Access

Synchronization

Inconsistent Outputs

Message State Task Partition Host

Netherite - Workflow

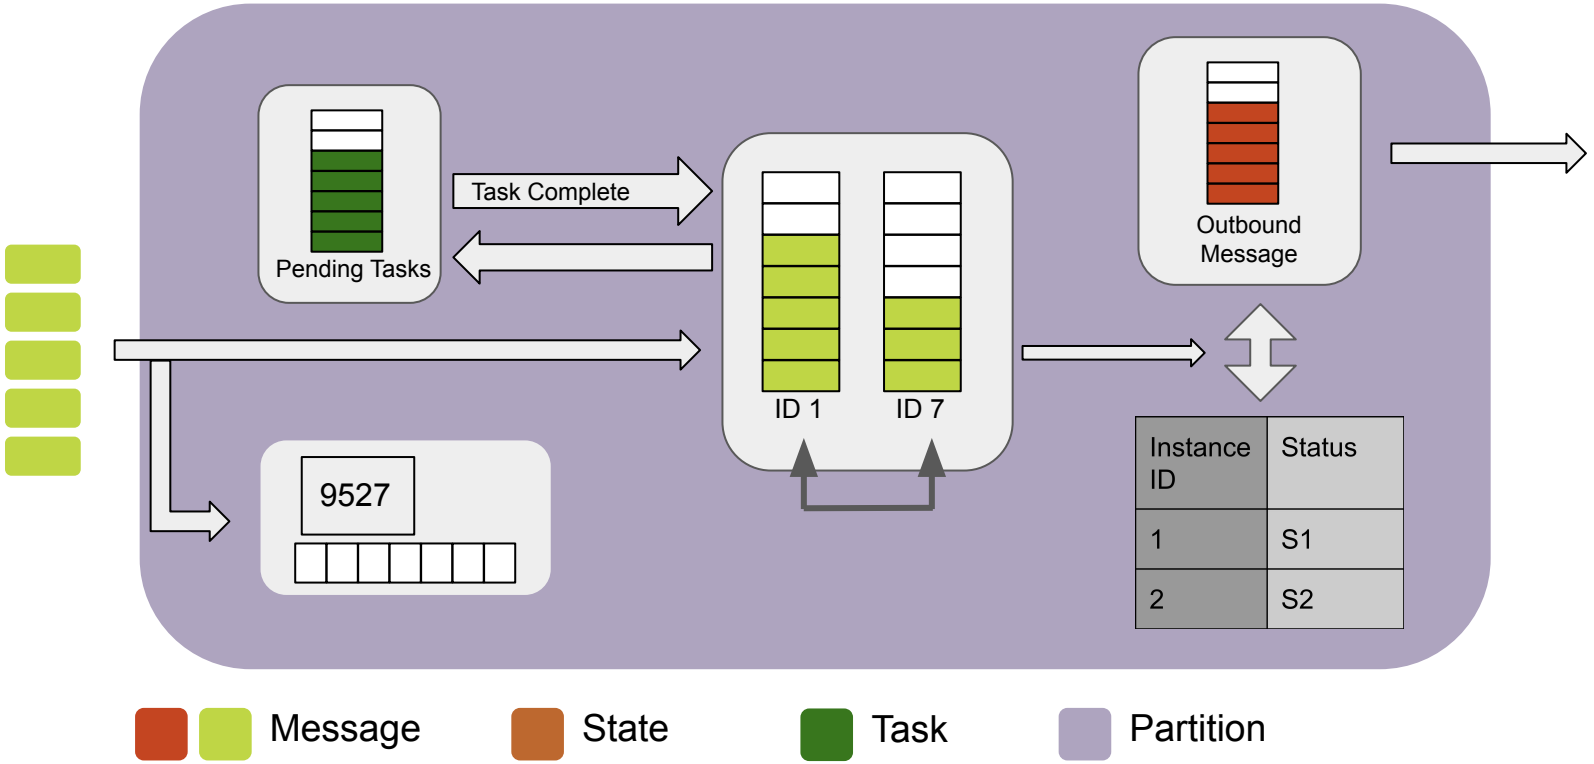

Netherite - Decoupled Persistence

Continuous log - **every** event

Occasional checkpointing

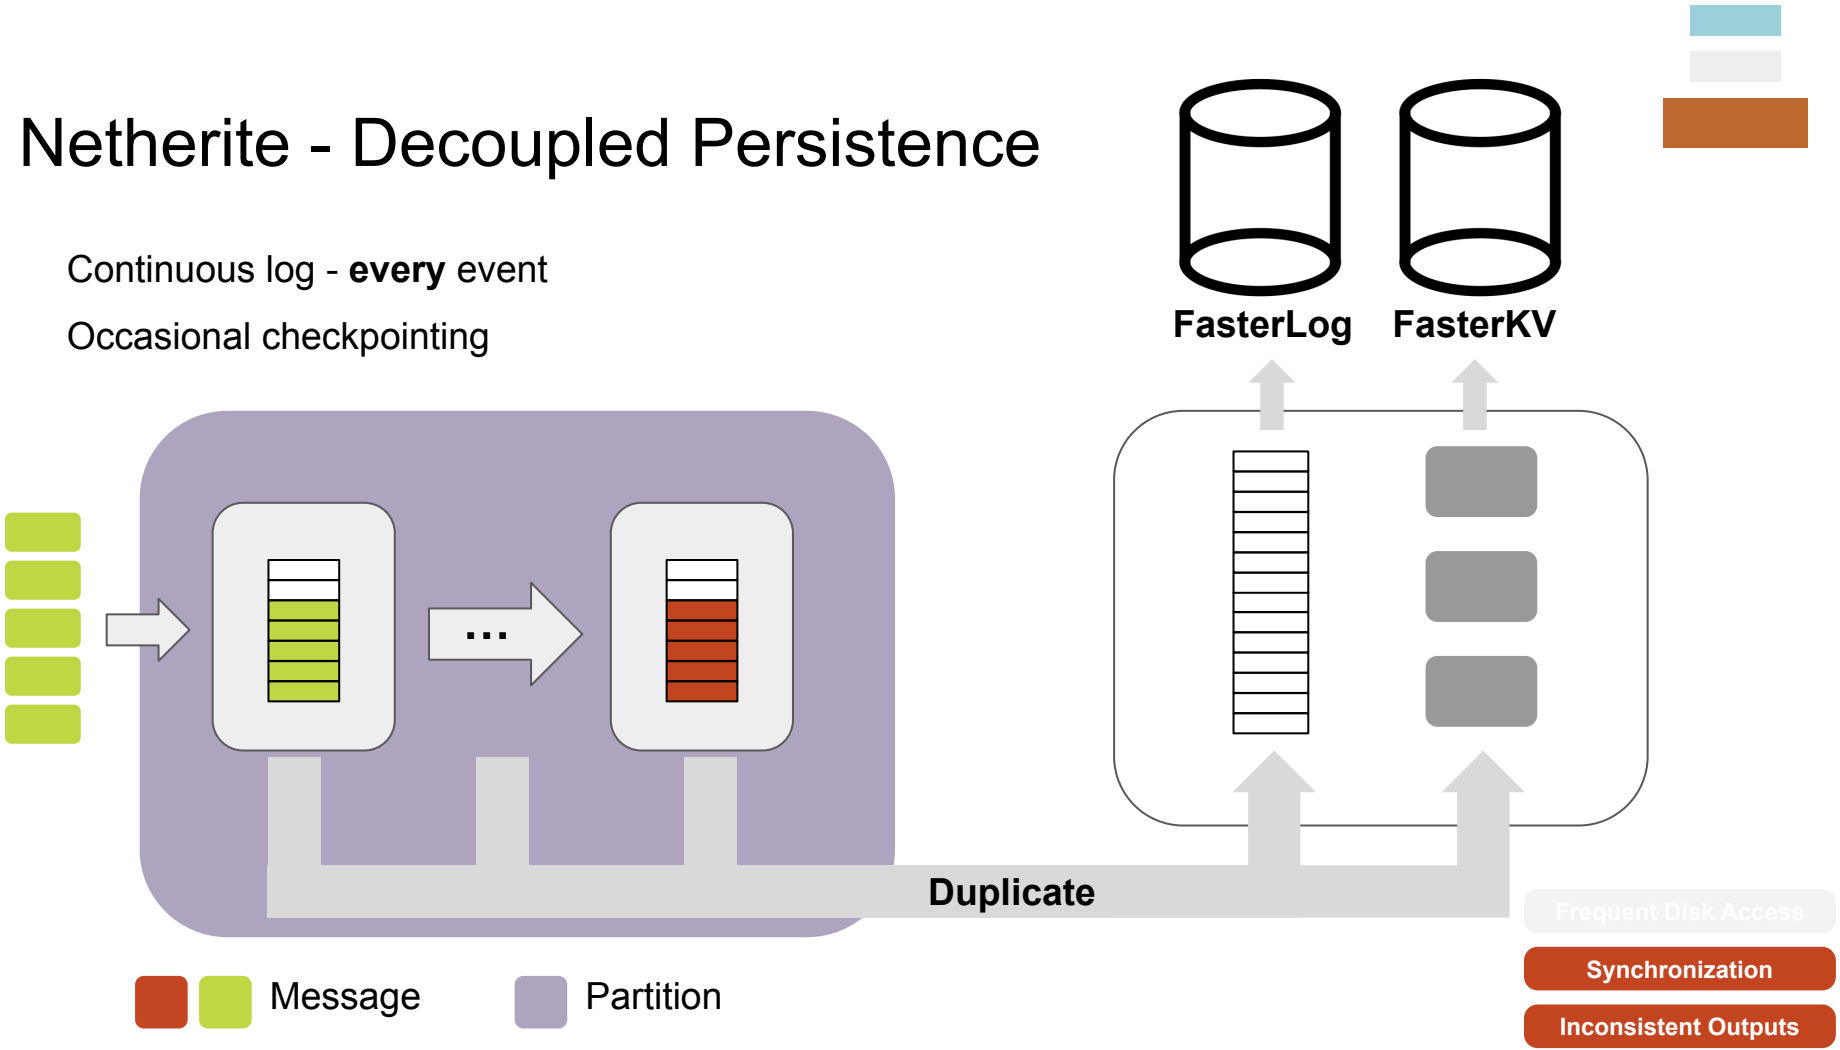

Netherite - Recovery

Frequent Disk Access

Synchronization

Inconsistent Outputs

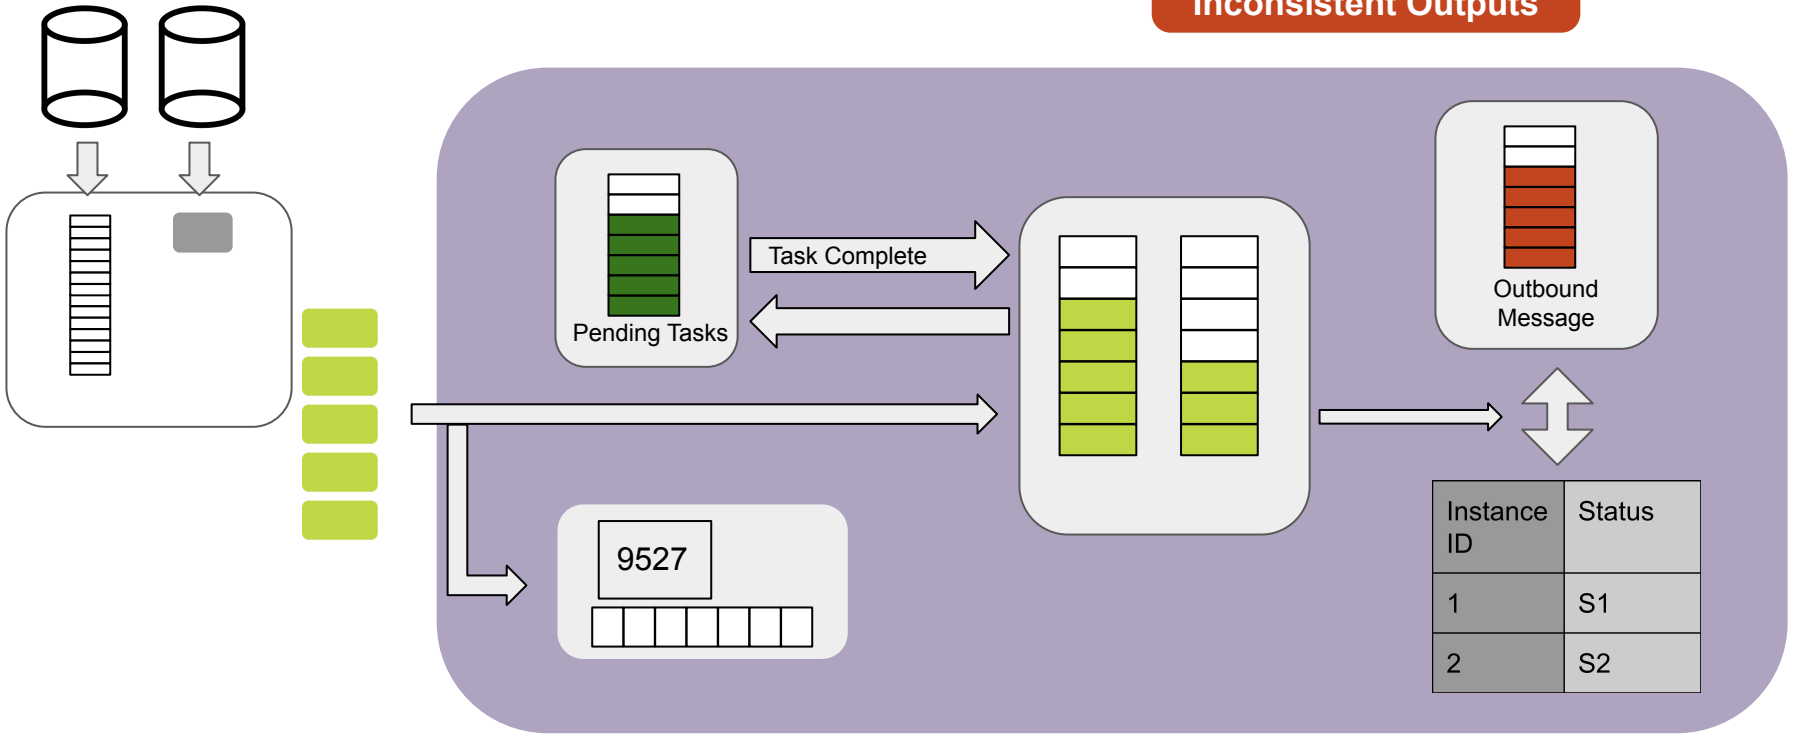

Partition State Persistence - Serialized Execution

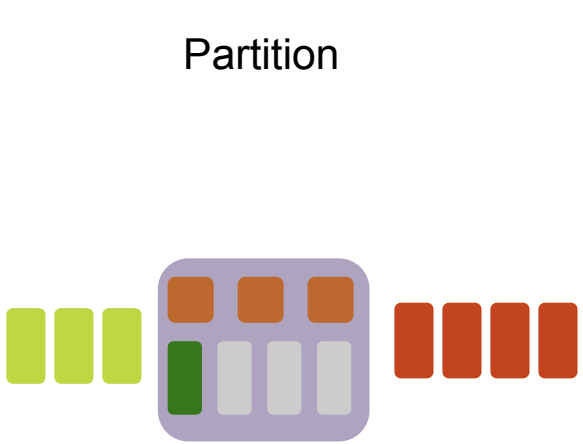

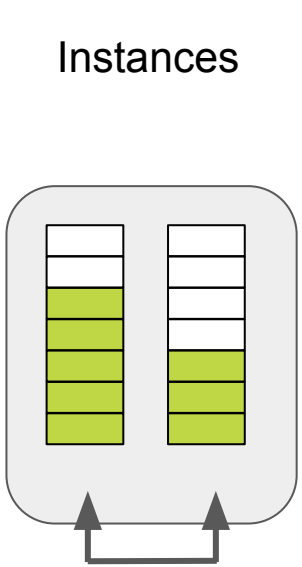

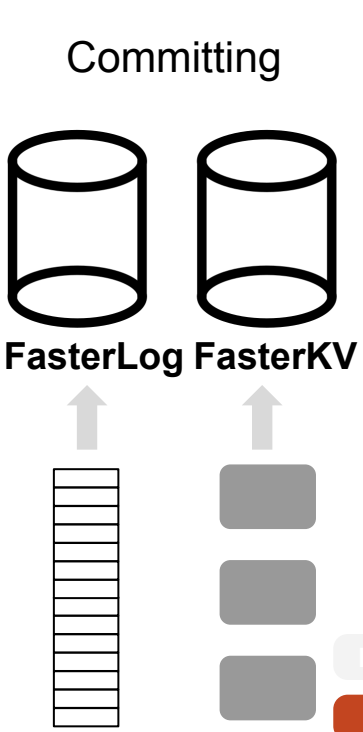

Message State Task Partition

Frequent Disk Access

Synchronization

Inconsistent Outputs

Partition State Persistence - Guaranteed Execution

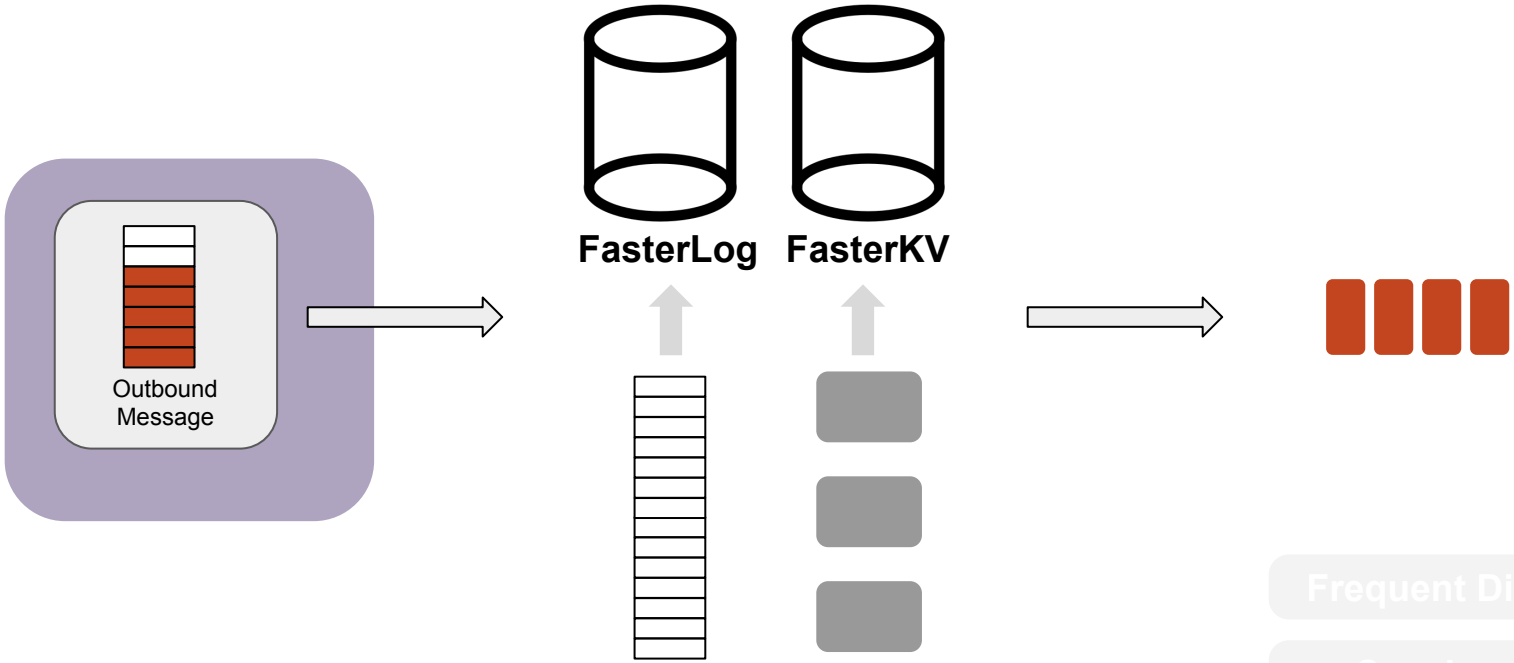

Frequent Disk Access

Synchronization

Inconsistent Outputs

Frequent Disk Access	Partition Level Commit
	Load Balancing
Synchronization Challenge	Serialized Execution
Inconsistent Outputs	Decoupled Persistence
	Recovery Mechanism
	Guaranteed Execution

Frequent Disk Access	Batch Operation	Partition Level Commit
		Load Balancing
Synchronization Challenge	Two-phase Lock Protocol	Serialized Execution
Inconsistent Outputs	Serialized Commit	Decoupled Persistence
	Partial History	Recovery Mechanism
	Rollback	Guaranteed Execution
	External Effects	

Evaluation

Evaluation …

Throughput

Storage traffic

Evaluation (continued)

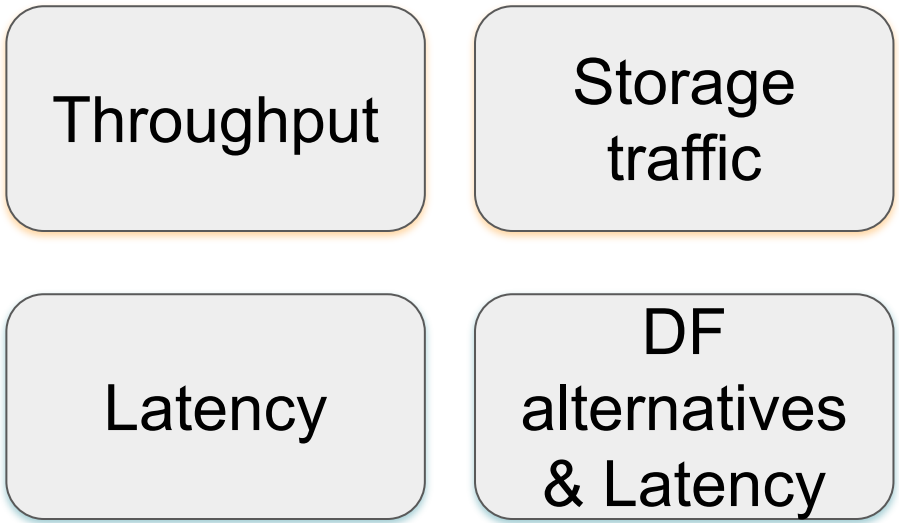

What is DF Alternatives?

Durable Function Alternatives

serverless functions with queues or triggers

Durable Function Alternatives

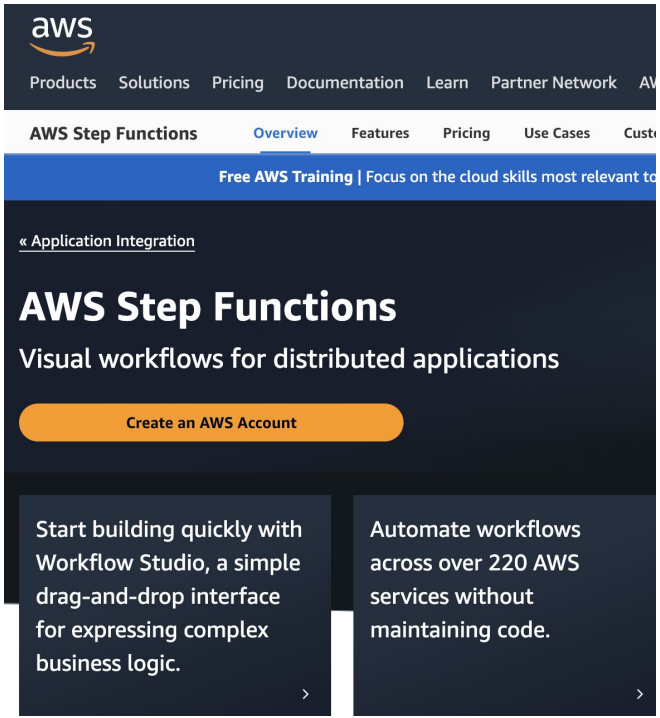

a workflow service

Use JSON

EVALUATION: Throughput Experiments

Hello5
Calls five talks in sequence

Bank
Reliable transfers between accounts

WordCount
Analyzes word frequency

CollisionSearch
Divide-and-conquer searches hash collisions

```
[FunctionName("Transfer")]
public static async Task<bool> Transfer(
    [OrchestrationTrigger] IDurableOrchestrationContext ctx)
{
  (string source, string dest, int amount) =
    ctx.GetInput<string,string,int>();
  EntityId sourceId = new EntityId("Account", source);
  EntityId destId = new EntityId("Account", dest);

  using (await ctx.LockAsync(sourceId, destId))

    nt bal = await ctx.CallEntityAsync<int>(sourceId, "Get");
    f (bal < amount)

      return false;

    lse

      await Task.WhenAll(
        ctx.CallEntityAsync(sourceId, "Modify", -amount),
        ctx.CallEntityAsync(destId, "Modify", +amount));
      return true;
} } }
```

EVALUATION: Methodology

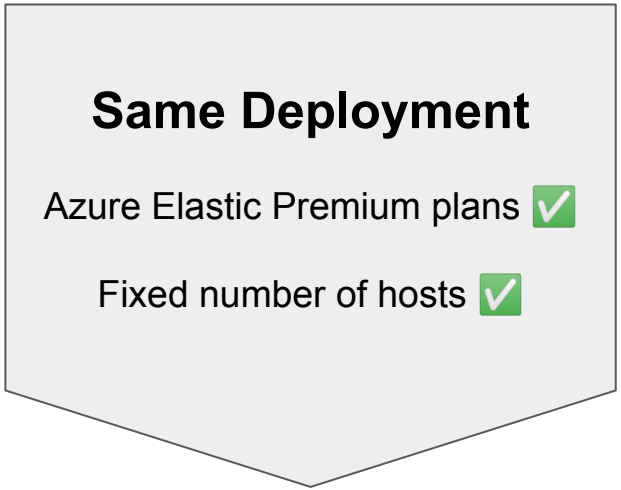

EVALUATION: Throughput Experiments

Does Netherite improve throughput compared to the original DF implementation?

EVALUATION: Throughput Experiments

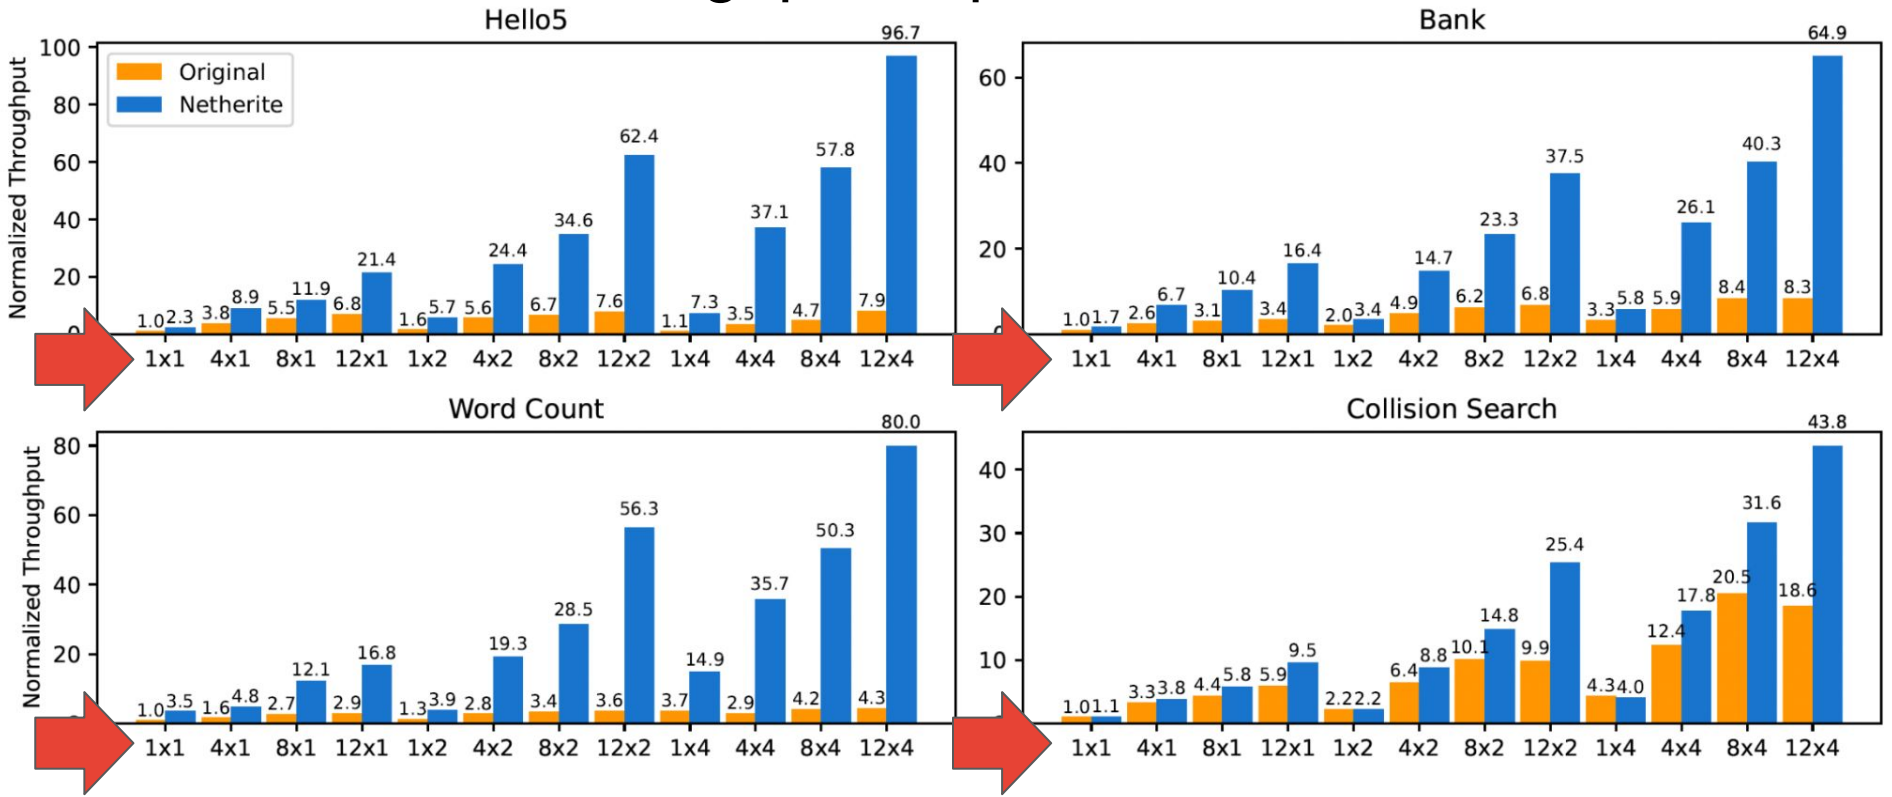

EVALUATION: Throughput Experiments

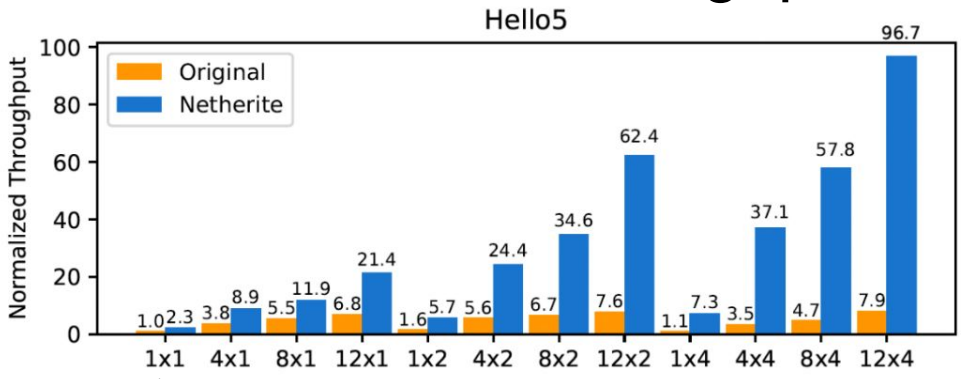

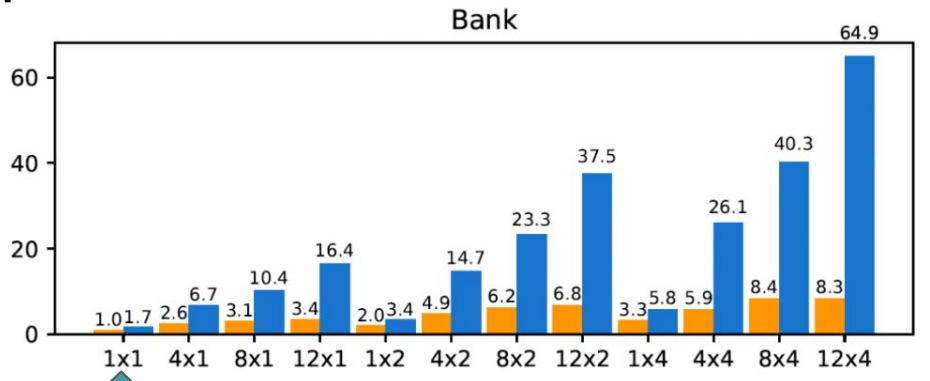

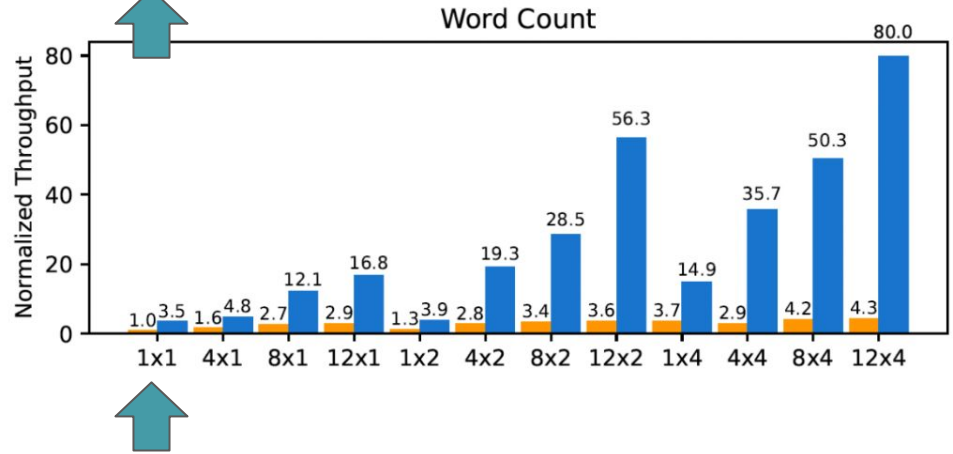

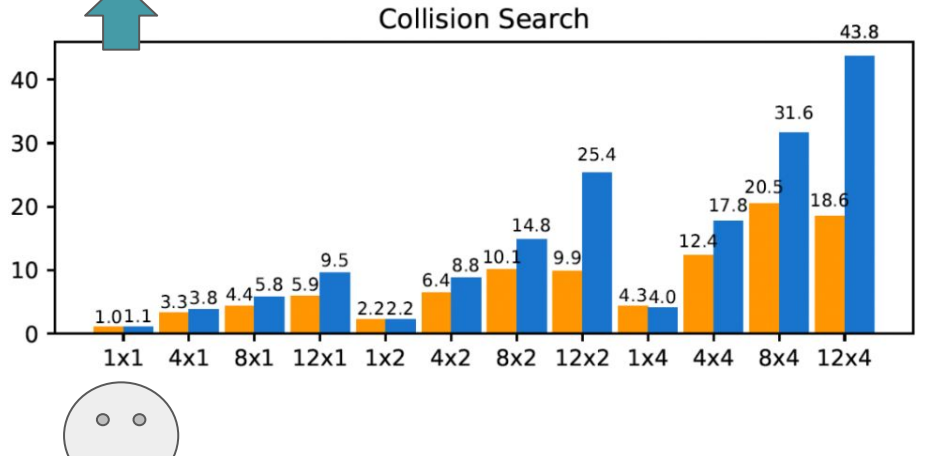

EVALUATION: Throughput Experiments

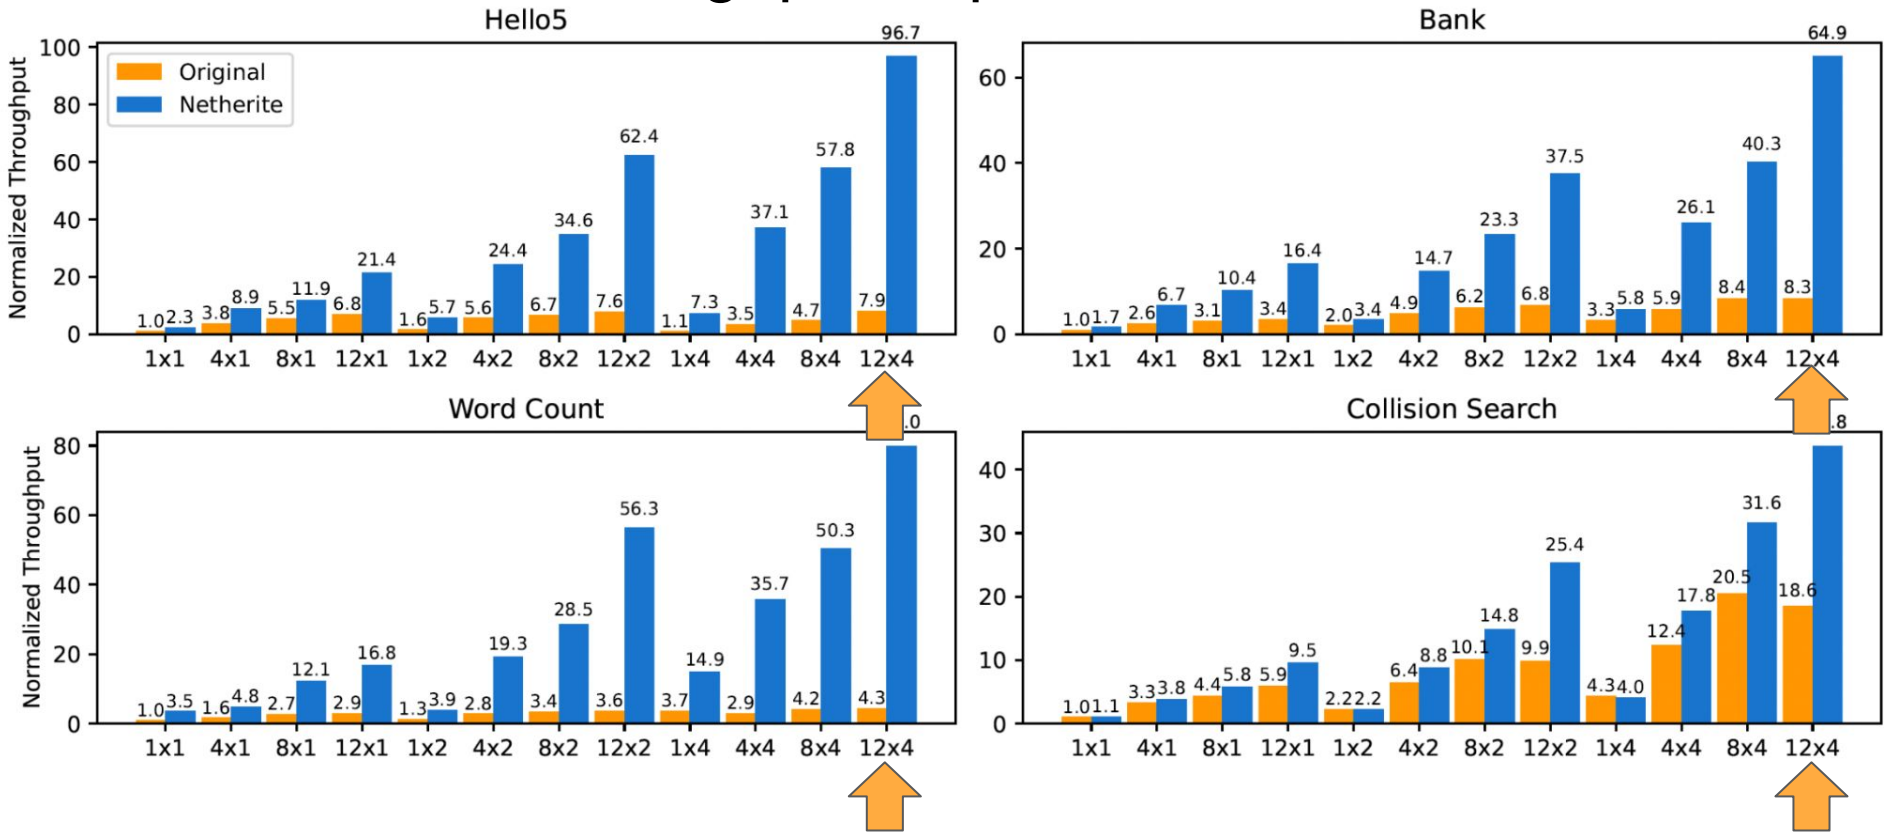

EVALUATION: Throughput Experiments

- Pipelining does not significantly affect the throughput.

EVALUATION: Throughput Experiments

- Pipelining does not significantly affect the throughput.

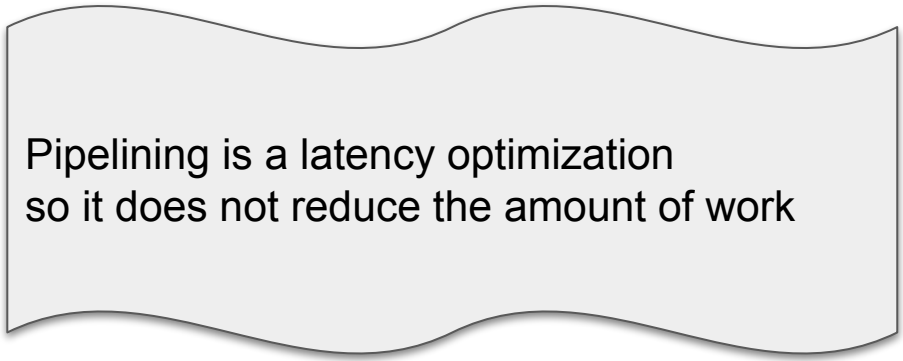

EVALUATION: Storage Traffic

Does Netherite reduce storage traffic compared to the original DF implementation?

The Number of Storage Request

Storage Access Volume

EVALUATION: Storage Traffic

Batching

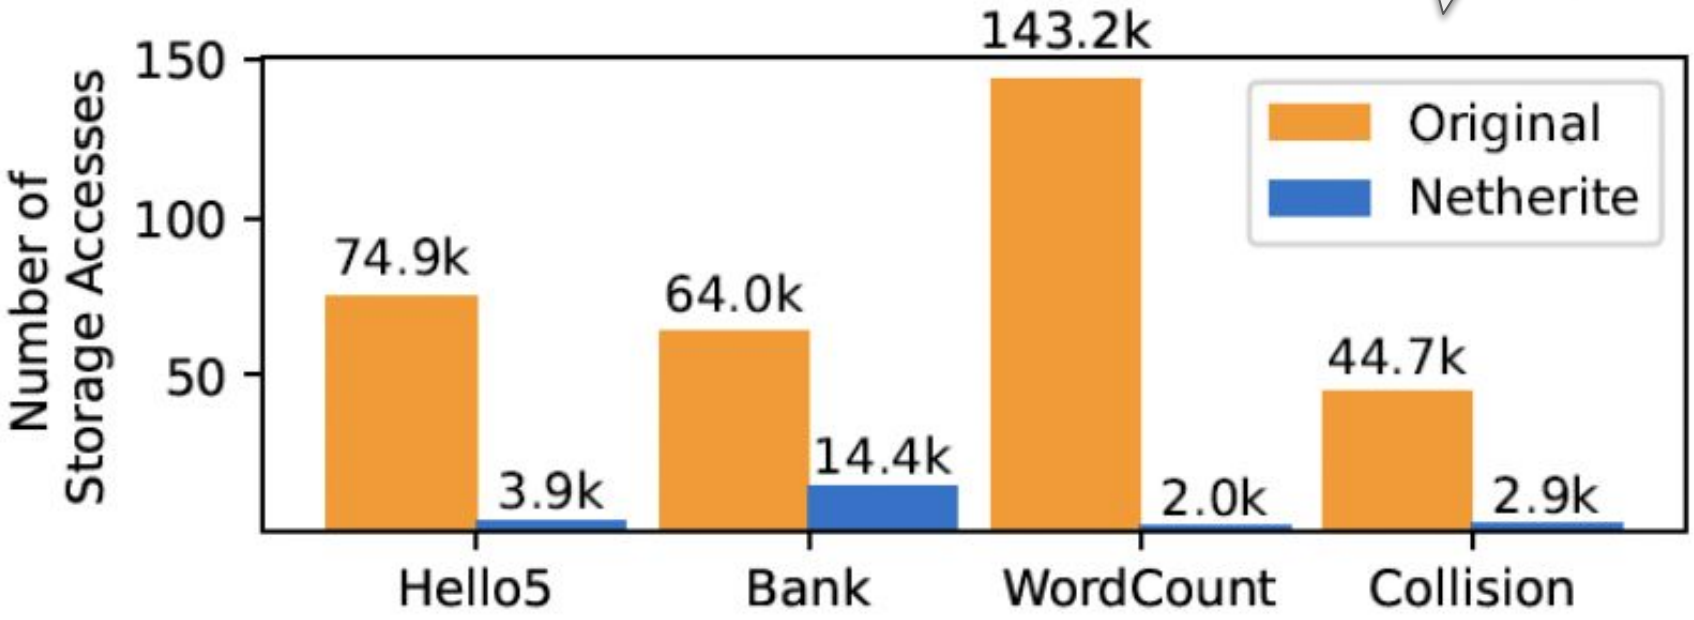

EVALUATION: Storage Traffic

Compact format

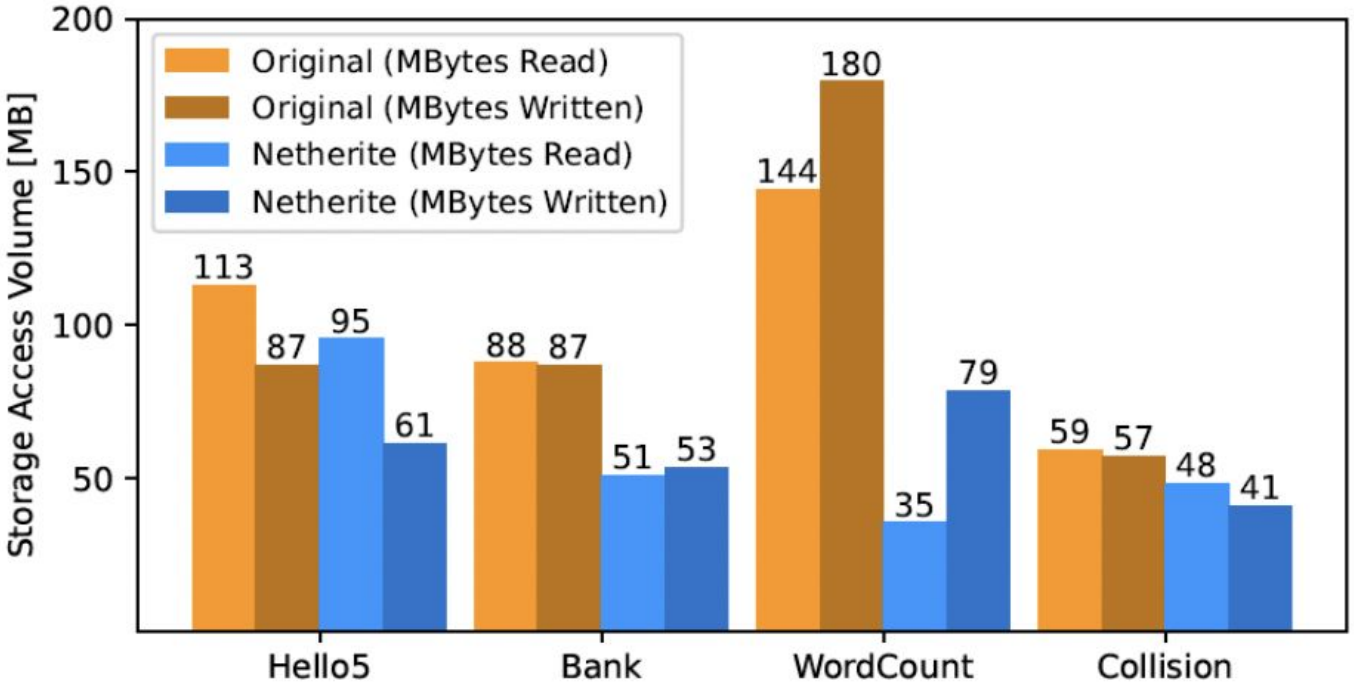

EVALUATION: Storage Traffic

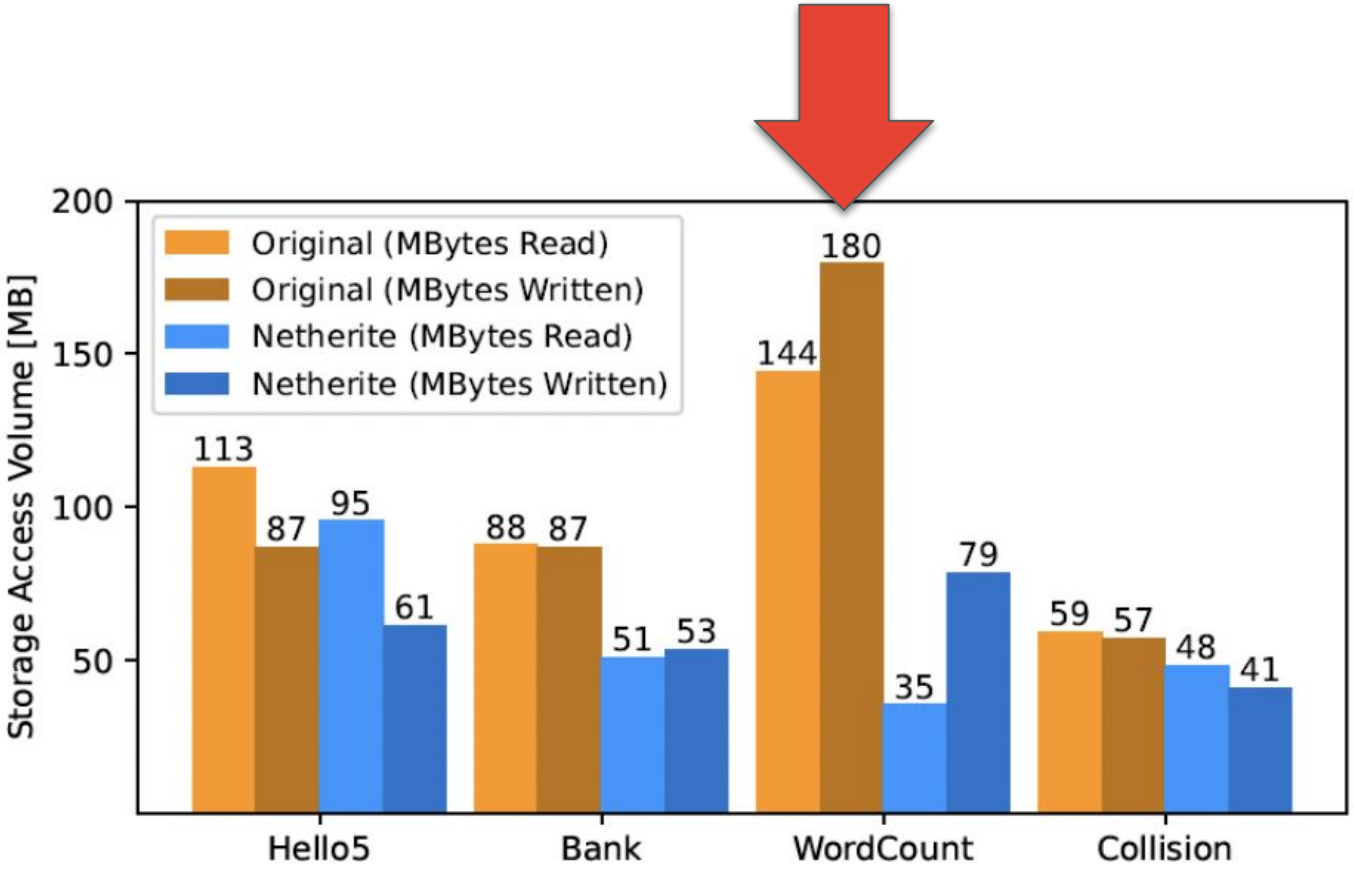

EVALUATION: DF alternatives

How does Netherite compare with DF alternatives when considering the latency of a single workflow?

1. Composition
2. Step Function
3. Original Durable Function

EVALUATION: DF alternatives

How does Netherite compare with DF alternatives when considering the latency of a single workflow?

Latency Only

EVALUATION: DF alternatives

Experiment with applications:

- Hello3
- Sequence
- Image Recognition

EVALUATION: Comparison to Common Alternatives

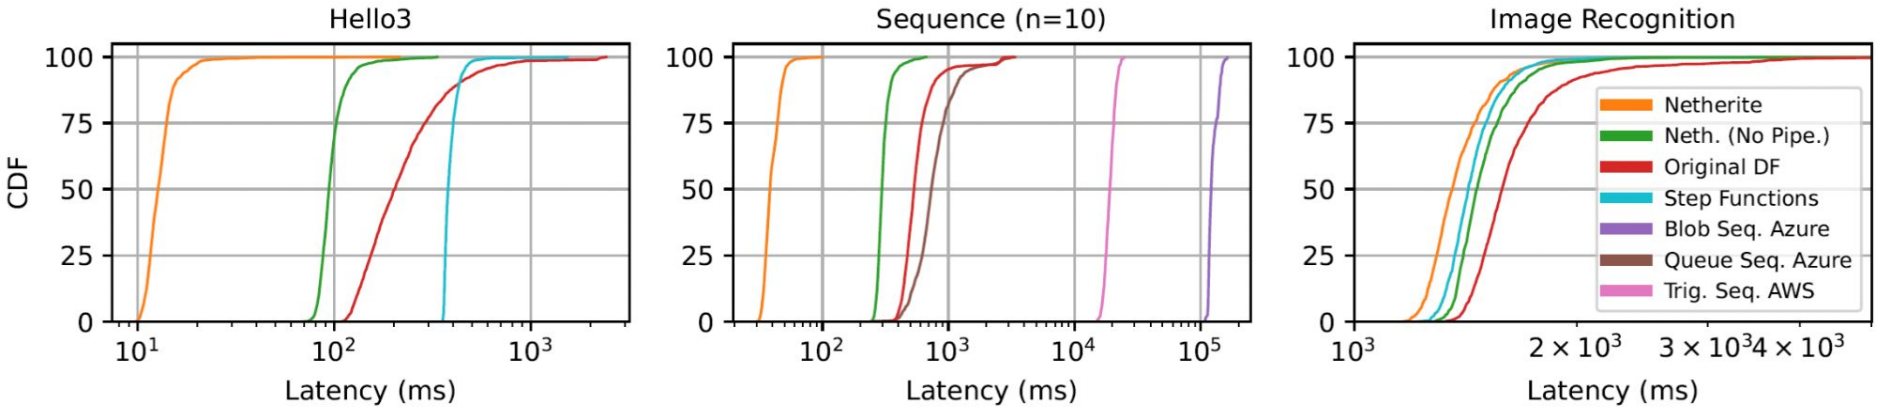

Conclusion

- Efficiency Improved …